

Free DVD
Slackware 13.0

LINUX



THE COMPLETE MAGAZINE ON OPEN SOURCE

VOLUME: 07 ISSUE: 8 October 2009 116 PAGES ISSUE# 81

ForYou

How To **Crack** Job Interviews

25 questions that are commonly asked.
Often they are not handled well.
Turn to page 22 for tips from a veteran on how to...

Published by EFY—ISO 9001:2000 Certified

India INR 100
US \$ 12
Singapore S\$ 9.5
Malaysia MYR 19



The Two Most Common
VPN Configurations

Setup OpenSolaris on the
Xen Hypervisor

Practices that Can Help You
Minimise Defects in Code

Try a Scheduler that's
Desktop-optimised

Limited Set of Constructs in C?
GCC Fills the Gap

Extend Web Applications with
WebHooks



Stop SPAM

Stop Virus

and...

Stop Paying!

Best
Welcome to the World's ~~First~~ Free Email Security Service

To activate SAFENTRIX for your domain, Visit <http://www.safentrix.com> now!

Tyrone[®]



GPU Optimized systems



2 TFLOPS in 1U
Up to 480+Cores

The Fastest 1U on the Planet
2 TFLOPS in 1U & 4 TFLOPS in Tower

4 TFLOPS in Tower
Up to 960+Cores
PSC: Personal Super Computer



- 2 Teraflops in 1U (240 cores per GPU)
- The Most highly optimized multi-GPU systems available
- Gold Level 93% efficiency 1400w power supplies
- Flexible system configurations with up to 4 GPUs

Powered by
SUPERMICRO[®]

Netweb
TECHNOLOGIES
ISO 9001:2000 Company

Corporate Office
B-1/ A25, 2nd Floor
Mohan Co-Operative Industrial Estate,
Mathura Road, New Delhi – 110044

New Delhi : +91 11 29942258/29942297
Bangalore : +91 80 41146565/40540000
Chennai : +91 44 43531009/26631975
Kolkatta : +91 33 23232055/23232056
Singapore : +65 64157239/92288638

Website : www.netwebindia.com
www.tyronesystems.com
Email : info@netwebindia.com
netweb@vsnl.com
mail@netweb.com.sg



25 questions

that are commonly asked.

Often they are not handled well.

Turn to page 22 for tips from a veteran on how to...

How To Crack Job Interviews

FOR YOU & ME

- 22 Cracking the IT Interview—FAQ
- 28 Slackware 13: Is it Really Worthy of the 21st Century?
- 31 Telepathy: Let There be Empathy!
- 34 ODF Olympiad 2009 Winners to Get Netbooks!
- 36 GIMP for Beginners, Part 4: The Selection Wand
- 40 What's Coming in Karmic Koala?
- 42 Liberating Young Scholars



BIZ USERS

- 48 You Can Bank On Open Source!
—A case study on IDBI Bank

ADMIN

- 56 Performance Tuning and Monitoring, Part 2
- 66 OpenSolaris on Xen, Part 1: The Ecosystem of the Xen Platform
- 69 OpenSolaris on Xen, Part 2: Exploring a dom0 OpenSolaris
- 76 Creating VPN with OpenVPN



GEEKS


- 54 Fine Tuning the Environment and Making Passes at the Command Line
- 72 The New Scheduler on the Block, Dedicated to Desktops



CONTENTS

DEVELOPERS

- 83 One Extra Step, Courtesy GCC
- 86 Defect Prevention and Identification
- 89 Python in Research: Computation with Polynomials
- 92 The Programmable Web
- 96 CPUID: The x86 Processor Identification Tool



COLUMNS

- 51 The Joy of Programming: The 'Struct Hack' Technique
- 62 CodeSport
- 98 A Voyage to the Kernel—Part 17: Segment: 3.6, Day 16


REGULAR FEATURES

- 06 Editorial
- 08 You Said It...
- 12 Q&A Section
- 14 FOSS Bytes
- 52 Tips & Tricks
- 104 Linux Jobs
- 109 FOSS Yellow Pages



LFY DVD: Slackware 13.0

This release brings with it many major changes, including a completely reworked collection of X packages, major upgrades to the desktop environments (KDE version 4.2.4 and Xfce version 4.6.1), a new .txz package format with much better compression, and other upgrades all around.



LFY CD: OpenOffice.org 3

+ POWERUSERS: •NetVault •Areca Backup •Bacula •FullSync •BackupPC
•Amanda •MySQLDumper; **DEVELOPERS:** •Gambas •gSOAP •Jython;
FUN STUFF: •Tremulous •Quadra; **UPDATES:** •Kernel 2.6.31

Note: All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported Licence a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence.

Editor

RAHUL CHOPRA

Editorial, Subscriptions & Advertising

DELHI (HQ)

D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26810602, 26810603
Fax: 26817563
E-mail: info@efyindia.com

BANGALORE

No. 9, 17th Main,
1st Cross, HAL II Stage,
Indiranagar, Bangalore 560008
Ph: (080) 25260023; Fax: 25260394
E-mail: efyblr@efyindia.com

CHENNAI

M. Nackeeran
DBS House, 31-A, Cathedral Garden Road
Near Palmgroove Hotel, Chennai 600034
Ph: 044-28275191; Mobile: 09962502404
E-mail: efychn2@efyindia.com

Customer Care

E-MAIL: support@efyindia.com

Back Issues

Kits 'n' Spares

D-88/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 32975879, 26371661-2
E-mail: kits@efyindia.com
Website: www.kitsnspares.com

Advertising

KOLKATA

D.C. Mehra
Ph: (033) 22294788
Telefax: 22650094
E-mail: efycal@efyindia.com
Mobile: 09432422932

MUMBAI

Flory D'Souza
Ph: (022) 24950047, 24928520; Fax: 24954278
E-mail: efymum@efyindia.com

PUNE

Zakir Shaikh
Mobile: 09372407753
E-mail: efypune@efyindia.com

HYDERABAD

P.S. Muralidharan
Ph: 09849962660
E-mail: efyhyd@efyindia.com

Exclusive News-stand Distributor (India)

INDIA BOOK HOUSE PVT LTD

Arch No. 30, below Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 022- 40497401, 40497402,
40497474, 40497479; Fax: 40497434
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Ratna Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2009. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

Editorial

Dear Readers,

It's the festival season, and is accompanied by a lot of good news. The IT biggies have started hiring again. Some of them have also announced their intention to restart the process of appraisals and increments.

But, did the demand for open source professionals ever go down? Our interaction with the industry indicated that while there was a dampening effect of the salary levels being offered, the demand for open source professionals remained the same as ever!

Whether it was the small and medium IT firms or the biggies—they kept looking for professionals with skills in open source projects. In fact, as more and more open source projects are being recognised as 'business ready', the demand for the skillsets to work on these projects has also been increasing. For example, projects like Drupal (a popular content management system), or Pentaho (a business intelligence software), are creating many jobs in India.

In fact, there's quite a demand for IT implementers from the Indian industry, as a whole. But the big issue is that everyone wants people with proven skills, and that's where the Catch-22 arises. If you don't get a chance to work, how will you prove yourself; and if you don't prove yourself, how will you get work?

Well, in the open source world, even this problem can be resolved. Pick up projects of your interest. Start reading about them, trying them out, and then contributing to them. If you do learn something, your knowledge will surely get recognised by the project team, and its community. That recognition is your 'experience certificate'. In fact, from all the success stories I have been hearing, there are very high chances that you will get a call for an interview, even without you having applied for a job!

Of course, preparing for that interview will be your next challenge. But, we have some valuable tips for you—as you might have noticed from our cover this month. S.G. Ganesh, who's been a regular contributor

to *LINUX For You*, has written a book titled 'Cracking the C, C++ and Java Interview', which has recently been published. He's picked some interesting chapters from his book to share exclusively with LFY readers. Turn to Page 22 to get valuable insights on how to handle job interviews, which have been shared by a veteran who's been on both sides of the table.

If you don't get a chance to work, how will you prove yourself; and if you don't prove yourself, how will you get work?

Here's a little secret: contributing to *LINUX For You* or *LinuxForU.com*, and getting your articles published can be the road to fame too. From authoring articles to authoring books, speaking at seminars, giving keynotes, and then perhaps becoming a top-notch consultant could be a great way to promote open source in India.

Wish you a joyous festive season on behalf of the entire *LINUX For You* team...

Best Wishes!



Rahul Chopra
Editor, LFY
rahul@efyindia.com

Linux Learning Centre

Pioneers in Training on Linux Technologies

Trained participants from over 38 Countries in 6 Continents

Linux OS Administration & Security Courses for Migration

LLC102: Essentials of Linux OS
LLC103: Linux System & Network Administration
LLC203: Linux Advanced Administration
LLC303: Linux System & Network Monitoring Tools
LLC403: Qmail Server Administration
LLC404: Postfix Server Administration
LLC405: Linux Firewall Solutions
LLC406: OpenLDAP Server Administration
LLC408: Samba Server Administration
LLC409: DNS Administration
LLC410: Nagios - System & Network Monitoring Software
LLC412: Apache & Secure Web Server Administration

Courses for Developers

LLC104: Linux Internals & Programming Essentials
LLC105: Programming with Qt
LLC106: Device Driver Programming on Linux
LLC107: Network Programming on Linux
LLC108: Bash Shell Scripting Essentials
LLC109: CVS on Linux
LLC204: MySQL on Linux
LLC205: Programming with PHP
LLC206: Programming with Perl
LLC207: Programming with Python
LLC208: PostgreSQL on Linux
LLC209: Joomla CMS
LLC501: Programming with OpenGL
LLC504: Linux on Embedded Systems

RHCE Certification Training

RH033: Red hat Linux Essentials
RH133: Red Hat Linux System Administration
RH253: Red Hat Linux Networking & Security Administration
RH300/301: Red Hat Rapid Track Certification Course

RHCSS / RHCA Certification Training

RHS333: Red Hat Enterprise Security: Network Services
RH423: Red Hat Enterprise Directory Services & Authentication
RH401: Red Hat Enterprise Deployment & Systems Management
RH436: Red Hat Enterprise Clustering & Storage Management
RH442: Red Hat Enterprise System Monitoring & Performance Tuning
RHS429: Red Hat Enterprise SELinux Policy Administration

RHCSS / RHCA Courses & Exams

RHS333: 3 & 12 Oct.; RH423: 10 Oct.

RH436: 06 October 2009

Exams: EX333, EX423 & EX429
on 5, 6 & 7 October respectively
EX401: 26 October 2009

Joomla CMS Course on 10 & 11 October 2009

RH301 from 12, 20 & 26 Oct. '09

RHCE Exam in Bangalore on 30 Oct. '09

LLC504: Linux on Embedded Systems 4 Day Fast Track Course starting on 10 October 2009

LLC410: Nagios System & Network Monitoring Software - Training from 10 Oct '09

Linux Support & Solutions

Installation, Setup and Support Solutions
for RedHat, Ubuntu, SUSE, CentOS Servers

For more info log on to:

www.linuxlearningcentre.com

Call: 9845057731 / 9343780054

Email: info@linuxlearningcentre.com

RED HAT Training Partner
RHCE & RHCSS Exam Centre



Linux Learning Centre Private Limited

Registered Office & Corporate Training Centre

635, 6th Main Road, (Adj.. Bank of India) Hanumanthnagar, Bangalore 560019

Tel: +91.80.22428538 / 26600839 / 26610999 TelFax: +91.80.26600839

Cell: 9845057731 / 9343780054 Email: info@linuxlearningcentre.com

www.linuxlearningcentre.com

LLC Satellite Centre - Bangalore

1291, 24th Cross, 30th Main,
BSK II Stage, Bangalore 560070

Tel: +91.80.26712928

You said it...



Hats off to the *LINUX For You* team for giving us an unimaginable distro. Yes, the Live Linux Gamers DVD that was bundled with the September issue proved that Linux is now an OS that's ready for gaming. I have a question... is it possible to install this distro; if so, please let us know as it would be very helpful.

—**Sam Benny, msambenny@gmail.com**

ED: A live installer for this distro would have been good, right? Unfortunately, there is no easy way to install the DVD. More details are available at: <http://live.linux-gamers.net/?s=faq>



The article on 'Develop a Simple Download Scheduler' was quite interesting, but what was suggested is difficult to implement daily. Hope someone develops a small GUI-based download scheduler that can be easily used by everybody. The requirement for a download scheduler is evident from the article and I can guarantee a million or more downloads in a month if implemented with a GUI. In fact, I was desperately in search of a download scheduler (most of my new Linux-using friends wanted one) when I came across this article. An appreciable effort—thanks to Sreekanth Balakrishnan. Still, I continue to hope that someone develops a front-end GUI for this.

—**E Sathiya Sekhar, sathiyasekhar@rediffmail.com**



I searched your whole site but I couldn't find which Linux distributions you are providing with the current issue. Please provide a detailed list of the software and Linux distros provided with the magazine on the *Current issue* page.

—**Anup Unawane, anup.unawane@gmail.com**

ED: Thanks for your feedback. We're in the process of conceptualising the 'Latest issue' Web page right inside *LinuxForU.com*—so that *lfymag.com* becomes redundant. We hope to implement it in October.



I need back editions from December 2008. Can I order them and pay by DD?

—**Anand, dreamscorner@yahoo.com**

ED: Due to a space crunch, we do not stock issues older than three months. However, our associates—Kits n Spares (www.kitsnspares.com) sell old issues that are typically up to a year old. Sometimes, they keep older issues too. To check if they have the issue(s) you want, shoot them an e-mail at kits@efyindia.com or call them at 011-32975879.



Why aren't searches based on author names working on *LinuxForU.com*? *Author Name* is one of the most common 'tags' that people use to start a search. Hope you will fix this.

—**Gurudutt Talgery, gtalgery@rediffmail.com**

ED: Thanks for reporting this bug. The problem was our search query had been searching for user names on the WordPress installation, rather than the first and the last names. However, we have managed to sort out the issue and hope it works fine for you now.

Hardware support—is Linux lacking?



First of all—congratulations on this very good website (and the *LFY* magazine, of course!).

But with Linux, there is a problem (actually, lots of problems!), which makes it impossible to run a PC without MS Windows. Why? Here's an example.

You can't print with USB printers. Maybe, it is working for some of us, but there are too many requests for help regarding this issue on the Internet. It is obviously not the printers or drivers, but CUPS and USB, or a combination of these, which creates the problem. You even read that some people switched back to Windows just because of the (non-)printing issue.

As I am not a developer, I can only guess what's wrong, but in the world of distros too, there are many flaws like this. I don't know whom but your magazine to address these issues to. I am sure your competence can do something in this regard, which will be useful to millions of otherwise frustrated users or wannabe users.

By the way, my system is 'home-made' with an ASUS M3M78-EM motherboard, an AMD 64-bit quad core, 4 GB RAM, an Epson Stylus Photo R230 printer, an Epson Perfection 3990 scanner and 2 hard discs. My operating systems are Mandriva 2009.1, Debian Lenny and unfortunately (yet unavoidably) XP Professional. It would be great if you can initiate a more cooperative system around the Linux core.

—**Hartmut Balke, hbpune@gmail.com**

Atanu Datta, assistant editor, LFY, replies: Thanks for your insightful feedback. I have one question w.r.t. the following hardware peripherals:

- ASUS M3M78-EM motherboard
- Epson Stylus Photo R230 printer
- Epson Perfection 3990 scanner

I'm sure these manufacturers provided you at least one driver CD with each of their products. However, in the absence of these driver CDs, do these devices work out-of-the-box on Windows XP Professional?

Get Certified With

Very

Talented

People



VENUS TECHNOLOGY PARK

Bangalore

Chennai

Pune



VENUS

No. 84, Rukmini Plaza,
1st Floor, DR. D.V.G. Road,
Basavanagudi,
Bangalore - 560 004,
Karnataka



vmware

authorized
TRAINING CENTER

Contact : 080-2650 0487 , 080-2662 2540
Mobile : +91-99013 01881 , +91-95661 25298
E-mail : aura@vtpinfo.net, vmware@vtpinfo.net
website : www.vtpinfo.net,

You said it...



Typically, whom would you call up or ask for support when any of these devices don't work as they should or as advertised in the absence of drivers? Is it the OS vendor Microsoft or the device manufacturers, ASUS and Epson? Would it not be the latter pair, rather than Microsoft? Similarly, calling Epson or ASUS for support would be more appropriate even with your Linux OS.

Linux, as of today, supports the most number of hardware devices out-of-the-box compared to Windows and other proprietary OSs. This is largely due to the occurrence of any of the following three scenarios:

1. The hardware manufacturer works with the Linux device driver developers—thus pushing its drivers to be included in the kernel, by default. For example, Linux is the first kernel to support USB3 specifications.
2. The hardware manufacturer makes the hardware specs of a device available—which enables the Linux device driver developers to write the required drivers.
3. The hardware manufacturer neither works with the Linux device driver developers, nor makes the device specs available—which forces the Linux developers to derive the driver software by reverse engineering.

There's one other mode of hardware support—the device manufacturer makes proprietary drivers easily available for different Linux distros to include, e.g., NVIDIA. This is arguably not the best of solutions (from the freedom point of view), but it works, nonetheless.

What's important to understand is the third scenario. Reverse engineering by guessing the specs of a device is not the easiest of jobs—yet many work at it just to make users' lives easier.

We, non-developers, could also help FOSS developers out—by asking the hardware vendors to support our OS. If they don't, then we should return their products saying that they don't work as advertised.

Better still, before buying a product, check if it works on Linux. Even when you know it works, still ask the sales team about it—they should know that Windows is not the only OS people use. And refuse to buy a product that doesn't support Linux.

It's small things like these that can help you, and all of us, out. Why? Because the moment the manufacturer realises that it's losing out on potential customers in the absence of Linux support, it has to rethink its strategy.

At the end of the day, manufacturers are rolling out products to make money by selling them. They support Windows only because they think that's their only market. We users need to make them realise that there are many people using Linux too. As long as we don't do our part, things won't change.

Hartmut replies: Following your suggestions I tried to

download the device drivers:

- ASUS, I just come to know, has no specific driver for any Linux system.
- Epson drivers are available and installed.

So, it is not a problem with any of the hardware drivers but with the handling of USB in the Linux system itself. Connecting, disconnecting or re-connecting at any of the USB jacks renders the entire USB system not to work, except for those units not disconnected or switched. Strange, isn't it?

Just check the loads of “usb-printer-problems” in the respective websites of all OS flavours and hardware brands—everywhere it's the same problem. The solutions recommended are as many as the questions. It is a Linux internal matter.

In a German website I found a question similar to mine, raised by a software(!) company. It was not able to solve the problem and switched back to MS because of the need to print.

Please find out what is wrong. I don't have the means to do this. Linux is using a lot of software made by ‘outsiders’, e.g., CUPS, which is an Apple product. It works with ghostwrite, foomatic, gutenprint It works, but who is going to standardise what is necessary?

You are right; we have to support and defend the freedom of our system. Salesmen want to sell—for them, ‘everything’ will work with Linux. Indeed, you will find a lot of ‘workarounds’ even for extreme hardware on the Internet. That shows the power of the Linux idea but also the ignorance of (at least some) people working for it. Sorry, for my openness but with Linux, it has to be so...

ED: *It's indeed strange that although you have the official Epson drivers, your printer still doesn't work. Here at EFY, we often connect HP printers using USB, to our systems, and they work fine. We would recommend you contact Epson, and also file a bug report about the problem in the bug tracking system of the distro(s) you're using.*

The FOSS ecosystem is a collaborative effort—that's why there are independent projects working on different software and applications. For example, CUPS has always been the print server software for Linux—Apple bought CUPS only in 2007. Anyway, we'd love to have the views of the community and our readers on the topic brought up by Hartmut. Thank you for your valuable feedback. Keep it coming! :-)

Please send your comments or suggestions to:

The Editor

LINUX FOR YOU Magazine

D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020

Phone: 011-26810601/02/03 **Fax:** 011-26817563

Email: lfyedit@efyindia.com, **Website:** www.linuxforu.com



Become an IT Expert

Join

USS

Ultramax System Specialist

100 % job Written Guarantee

Other courses & certification

SAS Certification
with exam preparation

Clinical Research
Live Project & Internship

RHCE & RHCSS

MCSE

CCNA

CCNP

Sun Solaris 10



Thane
6766 9228

Pune
4107 1111

Bengaluru
4211 5220

ULTRAMAXTM

INFONET EDUCATION

Dadar 6766 9222

enquiry.dadar@ultramaxit.com

[Www.Ultramaxit.com](http://www.Ultramaxit.com)



Q I have created a virtual server using OpenVZ, on which I have installed CentOS 5.3 and MySQL 5.0.45. After installing MySQL I was unable to start mysqld because of the appearance of the following error, according to the log file: Fatal error: Can't open and lock privilege tables: Table 'mysql.host' doesn't exist". Please help me configure MySQL.

—Suraj Ghugtia, Lucknow

To overcome this issue you need to remove the MySQL datadir. You don't have to worry about the data as you have a new installation. Check the location of your mysql datadir in /etc/my.cnf—by default, it is /var/lib/mysql—and remove it:

```
# rm -fr /path_to_datadir
```

Now go to /usr/bin/ (or whatever your MySQL path is) and run the following commands:

```
./mysql_install_db --user=mysql --ldata=/sqldata  
./mysqld_safe --datadir=/sqldata --user=mysql &
```

This should resolve the issue.

Q I was trying out the KDE 4.3-based openSUSE 11.1 implementation (LFY Sep'09 issue) and later installed it on a Dell Inspiron 1420 laptop. I would like to use the Reliance NetConnect USB on this laptop. However, I could not get any specific help on the Web or at forums.opensuse.org. I am hoping you will point me in the right direction. I have a Reliance Netconnect (CDMA 1x) ZTE / MG880 device. I was able to capture the dmesg output after I plugged-in and then unplugged the device.

```
usb 6-1: new full speed USB device using uhci_hcd and address 6  
usb 6-1: configuration #1 chosen from 1 choice
```

```
usb 6-1: New USB device found, idVendor=19d2, idProduct=fffd  
usb 6-1: New USB device strings: Mfr=1, Product=2, SerialNumber=3  
usb 6-1: Product: ZTE CDMA Tech  
usb 6-1: Manufacturer: ZTE, Incorporated  
usb 6-1: SerialNumber: Serial Number
```

On un-plugging the modem, dmesg shows the following:

```
usb 6-1: USB disconnect, address 6
```

—R V Mishra, rajyav@hotmail.com

To connect and use your MG880 modem, follow the steps given below.

Open a terminal and switch to being the root user. Now run a modprobe:

```
# modprobe usbserial vendor=0x19d2 product=0xfffd
```

You can find the product and vendor ID in your dmesg or by running the following command:


```
# cat /proc/bus/usb/devices
```

I assume that you have wvdial installed. Now just run:

```
# wvdialconf
```

After the successful completion of the command, check the /etc/wvdial file, which should have the following information:

```
Init1 = ATZ  
Init2 = ATQ0 V1 E1 S0=0 &C1 &D2 +FCLASS=0  
Modem Type = Analog Modem  
Phone = #777  
Username = <Put_your_phone_number_here>  
Password = <Put_your_phone_number_here>  
ISDN = 0  
SetVolume = 0  
FlowControl = Hardware (CRTSCTS)  
Modem = /dev/ttyUSB0  
Dial Command = ATDT  
Baud = 460800  
Stupid Mode = 1
```

Now run the wvdial command and wait for the modem to connect. **END** 

I am a
System Administrator / IT Manager
 I have done
RHCE / RHCSS / CCNA / MCSE !!
 However,
Still facing problems in the VoIP arena
 How can I get comfortable under
Asterisk Hood ?
 How can I improve my **VoIP skills ?**

Get enrolled into the fastest-growing next-generation
 IP Communication Industry worldwide?

Join Asterisk Training and Digium's
Official dCAP Certification at Enterux Solutions



Final Cut-Off Date for Registration

For Training:

31st October, 2009

Pay **Rs 1,25,000** + service tax and get 1 dCAP attempt **free**

For 1 dCAP Certification Attempt:

31st October, 2009

Pay **Rs 15,000** + service tax

*No more registration accepted after
 31st October 2009 for both training and certification*

Course Goodies Free with every Training

- 1 - Digium 1 Port PRI card without echo cancellation
- 1 - Digium 1 Port FXO + 1 Port FXS without echo cancellation Analog Card
- 1 - Polycom IP 330 IP Phone
- 1 - Training Courseware from Digium
- 1 - Asterisk Computer Laptop Backpack from Digium
- 1 - Asterisk T-Shirt
- 1 - Digium Calculator
- 1 - Digium Mousepad
- 1 - Orange ICE Digium Pen

Market Value Worth
Rs. 65,000/-
Free

Training Dates

9th - 13th November, 2009

Certification Date

13th November, 2009

Only Certifications

Upto 15th October, 2009

Pay **Rs 10,000** + service tax

Upto 24th October, 2009

Pay **Rs 12,500** + service tax



To Register:

Web: www.entvoice.com/training/signup.php

Phone: +912267339700

Email: sales@enterux.net



Athlon II X4 620 & 630—Quad-Core for Mainstream

With the recent introduction of Intel's new LGA 1156-based Core i5 and Core i7 processors, AMD faces even more pressure in the competitive upper-mainstream and high-end market segments. Phenom II is a great processor design, but it can only beat Intel's growing Nehalem family on price.

The new Athlon really isn't new, although AMD introduces two fresh core names for it: Propus (for the quad-core family) and Rana (for the triple-core).

Propus is a 2.6 GHz processor that's blessed with all the features of the Phenom II, including its 45 nm SOI manufacturing process and four cores with 512 KB L2 cache each. The chip also sports all of the extensions you'd want today: MMX, SSE, SSE2, SSE3, SSE4a, Enhanced 3DNow!, the NX bit feature (or execute disable, on Intel CPUs), 64-bit support, AMD-V virtualisation support, and Cool'n'Quiet to lower clock speeds and voltages during idle periods.

Since Propus is based on the Deneb design, all new Athlon II X3 and X4 processors can operate either on Socket AM2+ platforms with DDR2 memory or on Socket AM3 with DDR3. Clearly, the new processors represent an excellent upgrade option for older AM2 systems, especially if you consider the attractive \$100 price point.

Given this, the Athlon II X4 breaks with AMD's tradition of implementing shared cache memory in unified multi-core processor designs. The L3 omission is the main differentiator between the Phenom II and the Athlon II families, although there are obviously also clock speed differences. One thing about these two platforms that you always need to consider is whether the system is going to last. How easily is it upgradable? Well, with the AMD platform you can easily upgrade the processor to a Phenom II and the RAM is already DDR3.



Palm drops Windows Mobile for WebOS

Palm Inc has announced that it is going to stop making devices based on Microsoft's Windows Mobile, and will instead focus entirely on webOS-based smartphones like the Palm Pre.

For several years, this company followed a multi-OS strategy, offering models based on either Microsoft's Windows

Mobile or the Palm OS, which was recently replaced by the webOS. But no longer.

Jon Rubinstein, Palm's CEO, has announced, "Due to the importance of webOS to our overall strategy, we've made the decision to dedicate all future development resources to the evolution of webOS. Which means that going forward, our roadmap will include only Palm webOS-based devices."

Motoblur—the latest Android-based social media monster

In their bid to customise the open source operating system, Android and Motorola have come up with a completely different UI, christened Motoblur. It is interesting to note how much social media has affected today's mobile interface designs, living proof being the Nokia N97, HTC Hero and Motorola Cliq.

As you can see from the image alongside, the home screen looks like a very cluttered version of the Nokia N97 but has been heavily inspired by today's social networking needs. Some very interesting interface changes we found are:

- There are tons of social networking widgets. Be it FaceBook, MySpace, Twitter or Flickr, the OS has it all.
- The integration with the social networks is what blows your mind. One can check out updates of any contact, from all the social networks (supported) right in one single window. This had been a very useful feature of the BlackBerry OS 4.5+. I guess, it's time some one else took up the mantle too.
- All kinds of updates, notifications and messages have been streamlined into one single inbox, be it the SMSs, e-mails or social network updates.
- Getting a call? See the caller's latest Facebook/Twitter profile photo and Facebook status pop-up as the caller ID. Tap on a contact and see the person's entire social network history right

within the address book. It's a very comprehensive take on all your social networks.

- The phone (Motorola Cliq) is also equipped with a 5 MP camera, enough resolution for any social media enthusiast to keep sharing pictures within the network.



Want to be the Best ETHICAL HACKER & Security Expert?



After the grand success of the first edition that came out in June 2009
Leo Impact has come back with a 4 times **more powerful second edition.**

Security Expert
Average Salary
32 Lac/anum
Source: payscale.com

Even the most secure computers are Hackable...

- All E-mail addresses are hackable, including Gmail, Yahoo!, Rediff etc.
- Learn how to secure your systems and networks from hackers.
- All PCs can be hacked remotely using the latest tools and exploits, if they are on a network.
- Easily pass CEH (ver 6), CHFI, CISSP, CISA Certification.
- Learn Advanced Ethical Hacking:
 - Metasploit & Backtrack & Untraceable Hacking
 - Advanced Penetration Testing & Vulnerability assesement.

When you buy the book, you get:

- 2 DVD (18,500 tools, e-books, videos)
- E-mail Technical Support
- Free Lifetime Membership to Access Videos & Tools



Payment modes:

Credit Card, Netbanking, Cash Deposit & DD/Cheque

For more info. & online order: www.theseretofhacking.com

Order by phone: **+91. 9829944518, +91. 9953244518**, (10 AM to 10 PM)




LEO IMPACT SECURITY SERVICES PVT LTD

Corporate Office:

T8, Malyia apartment, near BJP office
c-schme, jaipur (Rajasthan) 302001
Email: contact@leoimpact.com

USA:

2029 Century Park East, 14th Floor,
California 90067 United States
Phone: +1 818 252 9090(USA)



Dell and Canonical introduce Ubuntu Moblin Remix

Canonical, partnering with Dell, recently unveiled Ubuntu Moblin Remix Developer Edition. The offering, based on Moblin version 2, is optimised for Dell's Inspiron Mini 10v netbook that became available since September 24th, this year.

According to a joint press release, Ubuntu Moblin Developer Edition “gives developer communities a great platform to develop and test Moblin and associated applications.” This piece of news mitigates potential concern about Moblin pushing aside Ubuntu as a popular mobile Linux operating system. Instead, the Dell-Canonical effort allows Moblin services and the Moblin graphical user interface to run atop Ubuntu.

However, Dell's promise to re-introduce desktop PCs with Ubuntu still remains unfulfilled.



Yahoo! may sell off Zimbra

As Yahoo works to reposition itself as a more consumer facing company, it is looking at some of the businesses it has acquired over the years and how they fit. One of them that doesn't appear to have a place in the new Yahoo is Zimbra, which, according to Channel Register [<http://ow.ly/qrs4>], is up for sale.

Yahoo acquired the open source e-mail company back in September 2007 for \$350 million. The goal was to use Zimbra's technology as a foothold into the white-label e-mail business and to improve Yahoo Mail.

Apparently, Comcast and Google are amongst the companies being courted to purchase Zimbra, although the sale price is likely to be far less than the \$350 million Yahoo paid. With Zimbra, Yahoo! had an opportunity to dominate the white label collaboration and e-mail space. Instead, it let the product languish and allowed Google to take a big chunk of the business.

Android's on a roll

Android seems to have caught the phone manufacturers' fancy, big time. This month, there have been a slew of releases on the platform. While HTC has been touting its brand new hero with HTC Sense, Samsung is also planning to introduce India to its very first phone with Google's OS in the core, called the Galaxy. Then there is LG Etna to drool on. Motorola has joined the bandwagon and designed a completely customised user interface for Motorola Cliq and Dext, all on top of the open source operating system, Android.

Moreover, to prove Android's popularity, Facebook has finally come out with an official Android app—something they had refused to do earlier, hoping the developer community would get its hands dirty. Google's Eric Chu has also announced several upgrades to the Android Market that will feature in Android 1.6. In a post on the official Android Developers' blog, Chu revealed that new changes to the Android Market would allow developers to add screenshots, promotional icons, and improved app descriptions to show off their apps and games better.



US FCC to mandate 'network neutrality' for the Web

The head of the United States Federal Communications Commission (FCC) plans to propose new rules that would prohibit Internet service providers from interfering with the free flow of information and certain applications over their networks. The guiding idea is that the big cable companies that build the information highway shouldn't exert much control over where users drive.

The FCC chief Julius Genachowski proposed two new guidelines, plus a strengthening of existing FCC standards:

- *Providers can't favour their own content*—Providers of broadband Internet services cannot discriminate against particular Internet content or applications. This would cover the emerging arena of wireless broadband services, delivered to smart phones.
- *Providers need to explain variable Internet speeds*—Broadband providers should disclose their network management practices to consumers. This might include policies that cause some users to face a squeeze on their connection speed during times of peak traffic on the network.
- *Providers can't limit access to lawful content*—These standards say that network operators cannot prevent users from accessing the lawful Internet content and applications.

Linux For You Magazine (Nov '09) brings

FOSS Training Special!

Because
Changing IT trends demand
more & more FOSS Experts!



Are you ready for it?

If you specialize in nurturing expert professionals in the field of

Programming (PHP, Python, Ruby, etc.) • System Administration • Networking and Security
• Application Servers (JBoss, GlassFish, etc) • Database Management Systems • Asterisk (VOIP services)

&

wish to reach out to Asia's richest community of

Developers, Administrators & End Users

Then join us in our quest to address the following concerns

→ Sudden boom in the demand for FOSS experts

→ Advantage of FOSS – Today & Tomorrow


→ The kind of expertise in demand

→ Success route to becoming an expert in
FOSS technologies

For any further information about the FOSS training special edition of LINUX For You, please contact

Dhiraj Khare +91 9811206582 or write to us at dhiraj.khare@efyindia.com





Linux is 'bloated and huge', says Linus

As the Linux kernel is becoming increasingly large and more complex, Linux founder Linus Torvalds says his job is getting easier. Speaking on a panel at the LinuxCon conference, Torvalds told the audience that the kernel development model is working better now than ever.



But Torvalds added that there are still areas for improvement and provided a very pointed comment about the current size of the Linux kernel. He added that whenever the kernel adds a new feature, the problem gets worse. That said, he didn't think that features are being added too fast and said that developers are finding bugs quickly.

Meanwhile, Linux 2.6.31 was also released in September (included in LFY CD), and you can read all about it at http://kernelnewbies.org/Linux_2_6_31.

Nero Linux 4: Reviving CD/DVD burning on Linux

Touted as one of the most prominent burning suites for Windows, Nero didn't stand back as far as Linux was concerned. Nero has recently updated its Linux offering and introduced a new version tagged as Nero 4.

The newer version improves overall integration and offers more pronounced support for optical drives and media handling in Linux. Version 4 introduces the Express Wizard, with which Nero simplifies the burning procedure in Linux, giving tough competition to its FOSS counterparts. The drawbacks with other players are incomplete burning integration and improper UDF implementation. Check out a video of Nero Linux 4 at <http://ow.ly/qg52>



ZOHO launches Discussions

Web-based productivity suite, Zoho, recently launched a brand new product called Zoho Discussions. Zoho lets any business, individual or organisation create public or private support forums where employees or customers can share comments around a particular discussion topic.

With the new product you can create a platform for discussion forums, similar to Google Groups. The differentiating factor is that your forums can be customised and branded to adopt the look and feel of your site. Zoho even lets you pick out a domain name that coincides with your site. Plus, Zoho Discussions can be integrated with many of Zoho's other productivity applications.

Similar to any forum, Zoho Discussion lets users create threads based on a particular topic. In terms of features, Zoho has focused on creating plenty of social tools to make the discussions more interactive and engaging. Aside from posting in the forums, users can interact in real-time through the built-in chat feature. Within the forum, users can create a profile, follow other users, bookmark particular threads and send private messages to administrators and users.

Zoho Discussion has a freemium model, with additional features like more storage, number of forums, number of moderators, etc, priced at \$25 and \$75 per month.

WebGL support awaits Firefox 3.7

It looks like support for the upcoming Web 3D graphics project WebGL is building up fast, as, just one week after it made its way into Webkit's latest experimental source-code releases, Mozilla Firefox 3.7 nightly-builds also sport the new feature. The project is still pretty much in the early stages and has a long way to go until standardisation, but the developers believe there should be a much sturdier implementation within several months.

The idea behind the project is to bring cross-platform, native 3D capabilities to browsers, independent of proprietary plug-ins, Adobe Flash and the likes, using standard Web technologies like JavaScript. It will use the <canvas> element from the upcoming HTML 5 standard and it has the support from most browser makers excluding Microsoft, of course.

One major step forward was when it finally became available, though not enabled by default, in the latest developer releases of Webkit, the open source HTML rendering engine initially developed by Apple and used by Safari, Google Chrome and several Linux and mobile browsers. And now it has also made its way into the Firefox branch codenamed Minefield, which will eventually develop into Firefox 3.7. Just like Webkit, it isn't enabled by default, as there still are several major issues to be handled, but curious users and developers can experiment with the new feature by going to the advanced-settings configuration page 'about:config' and enabling 'webgl.enabled_for_all_sites'.



IIT Madras
presents



Shaastra 2009
THE SPIRIT OF ENGINEERING

adapt. evolve. enjoy.

Sep 30th - Oct 4th




www.shaastra.org

Co-Sponsors



KAVIN™

Hospitality Partner



Gifts Sponsor FOSS Partner



mozilla



STATE BANK OF INDIA TAKING YOU FORWARD

ERICSSON



BOSCH
Invented for life




AIRBUS

Print Media Partner

EDUCATIONPLUS

Official Technical
Magazine Partner

ElectronicsForYou



Jolicloud: The new social OS for netbooks

Netbooks have become the buzz word today. As they become more popular, it becomes obvious that netbooks are very different from laptops and should be treated differently. Tariq Krim, the founder of Netvibes, is one of the many who are not satisfied with the OS running on their netbooks. And this led to the creation of Jolicloud (an OS based on Ubuntu Netbook Remix) that's designed specifically for netbooks, and the beta version of which is to be released very soon.

Jolicloud is currently at the closed alpha testing stage and only offers downloads on an invitation basis. The download is a LiveCD installer of around 600 MB, distributed in ISO format. You will have to use a USB creator software to mount the ISO file on a USB drive before it can be installed on a netbook.

One of the major highlights of the OS is its light weight and the absence of most of the applications one would normally use in your day-to-day life. Jolicloud has even decided to do away with the Synaptics Package Manager and replaced it with a Prism dashboard called My Jolicloud. The dashboard is the starting point of your Jolicloud Web-centric experience. It contains news and notifications of the software updates for your netbook.



The application panel functions just like the iPhone App Store where you can choose the applications you want and install it in your netbook, easily and quickly. You are not restricted to install only desktop-based applications. Most of the popular Web-based apps such as Google Docs, Zoho Office, and YouTube can also be installed in your system (basically, they are just a customised build of Prism).

The app also has a social network built in, where you can 'follow' other Jolicloud users (a la Twitter) and get updates on what they install on their systems (a la Wakoopa).

MontaVista 6 targets multiple platforms

MontaVista Software recently announced the general availability of new Market Specific Distributions (MSDs) of MontaVista Linux 6. By working closely with its semiconductor partners to deliver the right combination of features and functionality for each platform, MontaVista remains committed to providing the broadest hardware support in the embedded Linux market, and aligning the embedded Linux supply chain.

These MSDs are the first of a broad family of semiconductor feature-compatible, market-specific Linux distributions announced in May this year. Built on a common framework, and optimised for the respective hardware platform and its target market, an MSD is designed to support the full breadth of functionality provided by the semiconductor vendors, be feature compatible with the semiconductor vendors' Linux technology, and provide the value-add features and quality MontaVista is known for. MSDs may be customised and optimised for the target application, allowing developers to easily create a tailored software distribution that fully exploits the hardware specific features.

MontaVista Linux 6 was announced on May 12, 2009 and made generally available on August 14, 2009. More than 40 MSDs are being planned and developed and will be made available in the coming months. It also provides support for the broadest range of hardware platforms from each of the leading semiconductor vendors.

MS's open source strategist quits


Sam Ramji, Microsoft's key figure in engaging with open source developers, has left the building. He's quitting Microsoft and going to work for a start-up, leaving a number of Microsoft initiatives, including contributions to the Linux kernel and its 'CodePlex' open source code-sharing initiative, somewhat in the air.



That Microsoft has worked with the open source developer community at all would be news to many. Apart from being fodder for analysts and public relations, there's very little business incentive for Microsoft to really co-operate with free software in most of the areas it operates in. From Web servers to Web development platforms, compilers to dynamic languages, open source competes directly with Microsoft's own products.

Mesa 7.6 to be released before October '09

Mesa 7.6 provides support for a number of new OpenGL extensions with its software rasterizer and the Intel i965 driver mostly, there is the rewritten Radeon/R200/R300 driver that uses the buffer manager (a.k.a. the Radeon 3D driver re-write), GL_EXT_framebuffer_object support for ATI when using the TTM memory manager, proper OpenGL 1.5 support for the ATI R300 series (and VBO + OQ support), and the assembly shader rework. This, in fact, is just a small portion of the changes. Mesa 7.6 is quite a huge update compared to Mesa 7.5, which was just released in July.



WORLD SME CONFERENCE 2009

THE MOST ANTICIPATED EVENT IN THE SME ECO-SYSTEM IS BACK

Theme: 'My Way, My Vision' - From Giant Ideas to Giant Corporations.

At the World SME Conference, government bodies, SME associations and other stakeholders will outline their vision for the future of SMEs in India. Corporate India will share success stories on how companies were built from grassroots to global giants.

Eminent Speakers

"Grassroots to Global, a road map for Indian MSMEs - My Vision" By **Dr. A. P. J. Abdul Kalam, 11th President of India**, "The vision of the Ministry of MSME and the challenges in the near term" By **Sh. Dinsha J. Patel, Minister of State (Independent Charge), Ministry of MSME**, "Grassroots to Global - resurgence and the growth story", **Prof Anil Gupta, Executive Vice Chair, National Innovation Foundation**, "Initiatives of Ministry of MSME in 2009-10 in relation to 11th five year plan" By **Sh. Dinesh Rai, Secretary, Ministry of MSME**, "Preparedness of micro, small and medium enterprises to face the challenges in the globalized world" By **Dr. H. P. Kumar, CMD, NSIC**, "Tomorrow is Today" By **Mr. Sanjeev Bikhchandani, CEO Naukri.com**, Workshops on "IT solutions for MSMEs: Lip service or real solutions" By **Mr. Pradeep Gupta, Chairman, Cybermedia**, **Mr. Raman Roy, CMD, Quattro** and **Mr Vinnie Mehta, Executive Director, MAIT**, Workshop on "Alternative sources of funding for MSMEs" By **Mr. Saurabh Srivastava, Chairman Emeritus, TIE New Delhi**, "Indian innovations, global destinations - My Vision" By **Dr. Guriqbal Singh Jaiya, Director, Small and Medium-Sized Enterprises Division, World Intellectual Property Organization, Geneva**, "Collaborating with MSMEs from developed countries to enhance the competitiveness of MSMEs in developing countries" By **Mr. Bruno Valanzuolo, Chief Technical Advisor, Consolidated Project for SME Development in India**, Mentoring workshop on "Davids befriending Goliaths to create a win-win world" By **Mr. Nitin Nohria, Richard P. Chapman Professor of Business Administration, Harvard Business School**.

Registration fees Rs. 10,000/- (taxes extra), Bulk-booking discounts available. For details, log on to events.milagrow.in or call Neha 9958189838 or email at neha.bhatia@milagrow.in

Supporting Partner



Associate Partner



Media Partner



Print Partner




Print Partner



Magazine Partner





Cracking the IT Interview —FAQ

This introductory chapter from the book 'Cracking the C, C++ and Java Interview' is in the form of Frequently Asked Questions (FAQ). It answers many of the basic doubts most students and young programmers have. These FAQ also clear a few misconceptions about attending interviews.

General questions

1 What skills do IT companies look for in prospective candidates?

IT companies look out for various technical skills and soft skills in candidates. In soft skills, communication skills are very important. Other soft skills include presentation skills, team-work, writing skills, etc. In technical skills, companies expect the candidates to have good expertise in their area of graduation. For computer science students, evergreen technical skills are C, Unix, operating systems and networking. Candidates with these skills are likely to have better chances of getting a job.

2 When should I start preparing for placements?

For soft skills, it is better to start preparing one year before the placements start. For technical skills, it is beneficial to focus from the first year itself. Otherwise, you can start revising important subjects—from the placements point of view—from one year before placements.

3 What is the difference between a CV and a résumé?

A CV (Curriculum Vitae) is a document prepared by a student or a fresher searching for a job. It provides the academic details of the student. A résumé is prepared by a person having work experience, which emphasises job history and on-the-job skills and experience.

4

What should or should not be there in my CV?

A good CV will have the following sections: objective, personal strengths, academic background, academic achievements (if any), extracurricular activities, project details, areas of interest, and personal and contact details. It should have neat and simple formatting. Ideally, it should be of 2–3 pages.

Some characteristics of a bad CV can be as follows: overly complex formatting or styles; typos, grammatical mistakes; either too short (1 page), or too long (> 3 pages); too many personal details (e.g., parent's occupation, 3rd consolation prize in Rangoli competition); list of weaknesses, etc.

5

How detailed should the CV or résumé be?

The CV or résumé should ideally have adequate details about one's job history or academic details, project details and achievements.

The job history should be from the current work experience to details of the first job (i.e., in reverse chronological order) giving the name of the company, role/position and number of years (preferably with joining and leaving dates). The academic details should cover the college/university, course, year of passing, percentages/CGPA, etc. Project details should include the title, where it was done, software used, and a short description of the project. Other sections can be very brief and to-the-point.

6

How should I prepare for written tests?

To prepare for written tests, the following should be taken care of:

- Know the general format of the question papers from IT companies. Typically, most of the IT companies cover some or all of the following topics:
 - programming aptitude (C, data structures, algorithms, etc)
 - quantitative aptitude
 - analytical ability
 - reasoning (logical, critical, etc)
 - verbal skills (synonyms, grammar, composition, etc)
 - puzzles
- Try solving the previous years' sample question papers
- Attempting to solve previous years' sample question papers is very important. This helps you know where you stand, get experience in answering questions quickly in the actual written test, and, in general, become confident of clearing the written test.
- Read relevant books
 - GRE (Barron's guide)
 - R.S. Aggarwal's aptitude books (quantitative aptitude, reasoning, etc)
 - Shakuntala Devi's puzzle books ('*Puzzles to Puzzle You*', '*More Puzzles to Puzzle You*', etc)
 - Technical books ('*Let us C*', etc.)

7

How should I prepare for attending an interview?

In the first place, know the basic details about the company—its main business, size, etc. If possible, visit the company's website and get to know the general details about the company. This helps to show that you are interested in the company when related questions are asked in the interview.

An experienced person looking for a job change needs to know about the current position he is applying for (the job profile), what is expected from a person in that position, and how he can fit into that position.

An important aspect in cracking the interview is your attitude (i.e., how you present yourself). Show keen interest, be attentive and listen to the interviewer. Other aspects to look out for are eye contact, body language, appearance, way of speaking, showing respect, etc. Also, be on time for the interview (e.g., start early if traffic jams are common in your city).

Knowing the latest advances in technology and other happenings in your technical domain would be an added advantage.

8

How many interviews do I have to clear to get a job?

It is usually a minimum of two interviews: a technical interview and an HR interview. In some cases (e.g., if the company you've applied to is in another city or country), a telephonic interview is done to screen the candidates before calling them for face-to-face interviews. In major Indian IT companies and MNCs that are product-based, there will be multiple technical interviews. Unless the hiring company is satisfied with your technical skills (and communication skills) and you clear the HR interview, you cannot get the job. Your social status makes no impact in the interviews.

11

I am a class topper. Should I attend only selective companies in our campus interviews because I am sure to get a job?

You're overconfident and this can spoil your chances of getting a job. Being a class topper obviously gives you a better chance of getting a job. But remember, the skills required for getting a job are different from getting high marks. It is better to get a job first and then start being selective about your 'dream company', which you may want to join as a second job. Also, if you're selective, you lose opportunities to attempt written tests and interviews and hence when your 'dream company' comes for placement, you'll be ill-prepared for it.

9

Why do companies have an HR interview in addition to the technical interview(s)?

The HR interview has two objectives—(1) to check if you're fit for the organisation; and (2) to check if your requirements match that of the organisation. If you have good communication skills, a friendly personality, a positive attitude, and keen interest in learning and contributing, and in addition, if you're an effective team player and it is likely that you'll stay for at least a few years, it's most likely that you will be beneficial to the organisation. An HR interview is intended to check all these aspects. In other words, an HR interview checks your soft skills, attitude, and if you're the 'right fit' for the company.

12

I got a job, but I did not get a call to join the company. Should I keep on waiting indefinitely?

Don't worry. If the job environment is bad, it is natural that your joining date can be very late (in some cases, the delay can be more than a year!). Keep in touch with the HR department of the concerned company and get to know the status from other candidates who have got placed in that company. Explore other alternatives: search for another job (who knows, you might be destined for a better job!); join some job-oriented courses; do certifications in the areas of your interest

10

Why is an HR interview important?

Simple—only if you clear this interview, you'll get a job! Also, if you get the job, the HR department (of course, after consulting with the manager of the team for which you're recruited) decides the pay you'll get, your roles and responsibilities, etc.

13

That brings me to another question. Do certifications help in getting a job?

Yes. Today certifications are an effective way to demonstrate your expertise in a particular technical domain. Getting relevant, valuable certifications can significantly improve your chances of getting a well-paid job.

THE COMPLETE MAGAZINE ON OPEN SOURCE

LINUX BETA
ForYou

Your favourite Linux Magazine is
now on the Web, too.

LinuxForU.com

Follow us on Twitter @LinuxForYou

14

I have been searching for a job for more than a year and I still have not got a job. What are the options that I have?

Don't lose hope. There are examples of candidates who have good jobs and successful careers, who couldn't get any job initially. Continue to search for jobs, but also explore other options in this situation:

1. Go in for higher studies and improve your academic qualifications.
2. If you have a bachelor's degree, consider enrolling for a master's degree such as an MBA, M.Tech, or other courses, depending on your interest.
3. Consider joining advanced courses, such as a post-graduate diploma programme from C-DAC.
4. Do some certification courses in your areas of interest, which can improve your chances of getting a good job.
5. Join some evergreen job-oriented courses such as software testing and technical writing.
6. Depending on your interest, take some specialised courses such as advanced animation, or CAD-CAM, from established institutes like Aptech, NIIT, etc.
7. Learn any new computer skills: new programming languages like C, C++, Java, operating systems (Unix, Linux), applications (Tally, etc). Such job skills significantly increase your chances of getting a job.
8. Network with your college seniors, relatives or friends who are already working and ask them to forward your resume to their HR departments.
9. Send your CV to companies both directly and through consultants.
10. Improve soft skills: communication skills, presentation skills, learning foreign languages (Japanese, French, German, etc).
11. Do software projects: it adds value to your CV. Try doing a project from a reputed organisation (MNCs, PSUs, government organisations, etc). Don't pay for doing projects. Rather, try working as a trainee without getting paid or with a minimum stipend (to gain experience).

15

Is it necessary to change jobs frequently to get a better pay and position?

No, it is not a good idea to change jobs frequently. Try to stick to a company and work there for at least 3 to 5 years.

There are good reasons why one would change a job for professional reasons (better pay, career advancement, new work environment, new area of work, overseas work assignments, etc), or personal reasons (getting married, want to live with parents, etc). It is perfectly acceptable to change a job for such reasons. However, don't change your job frequently. There are many reasons why we should avoid 'job hopping'.

Potential employers look at the job history of candidates before selecting them. If a person has changed jobs often (say, 5 jobs in 5 years!), it is very likely that the person will do so in future as well, so employers prefer not recruiting such candidates. Typically, it requires around 6 months to become productive in a new organisation and start contributing. If you leave the job within a short period—say within a year—it is a loss to the company because of many reasons: the company has ramped you in your new job and that effort is lost, the company has to spend again to recruit a new person for your position, the work gets pending till the time the new person on board becomes productive, etc. So it becomes difficult to get a new job if you are a 'job hopper'.

It is better to take a long-term view about your career. It takes at least 5 years to learn enough about the job, the company, the technology, become highly productive, and make significant contributions to the company. The pay and position we get in the company depends on the level of contribution we make to the company. If you find 'your kind of job and company' and stick to it, and focus on contributing to the company, you'll naturally grow and earn better than if you keep shifting jobs. 'Focus on learning than earning'—that's what all the successful people have done!

Some FAQ to candidates and how to answer them

16

Tell me about yourself.

This is an open-ended question that interviewers ask at the beginning of the interview to know more about you. They also use this question to get an idea of how you look at yourself and your achievements.

Briefly explain your professional background, the projects you've done, significant contributions you've made in your previous jobs, and conclude with a note about your personal background and a few points on your positive personal characteristics. Don't talk for an hour; make it short and to-the-point. Also, don't overemphasise your personal details.

17

What are your strengths and weaknesses?

This is a question asked to check how you look at yourself and also how your strengths can contribute to the team.

Be honest and tell what you consider as your strengths ('I learn new skills fast', 'I am an effective team player', 'I have good leadership skills', etc). Provide supporting details for your strengths ('I learn new skills fast. In my previous job, I had to learn scripting. I started to do shell programming from the next day itself, and I did it!'). For weaknesses, don't elaborate too much; some weaknesses can cost you your job ('I can't resist stealing if I see costly mobiles!').

20

Why are you planning to leave your current job?

Be careful in answering this question.

Usually acceptable answers are: 'Looking for better pay'; 'looking for a better role and growth opportunities'; 'Got married and had to shift to this city.' Bad answers: 'I didn't like my old boss!' (you're too frank to get this job!); 'The project is nearing the deadline and I don't want to work in that hectic schedule' (you can't desert your project when it is in a critical situation!); 'I had to work!' (come on, you're paid for doing your work!).

18

What do you know about our company?

This question is to check if you would be interested in working in that company (if you're a kernel hacker, it is unlikely that you'll be interested in Web programming, assuming that the company develops Web-based software). It is also to check if you're keenly interested in joining the company—if you're going for higher studies and attending interviews 'just-like-that', then you would not have shown much interest in knowing more about that company, right?

To answer this question you should prepare before attending the interview. Visit the company's website to know about the company. If you know anyone working in the company, contact them and get an idea of what the company works on, which countries (or states) it has a presence in, the kind of projects or products they are working on, etc. An overview of the company is more than enough.

21

Tell us about some challenges you faced in your previous job and how you overcame them.

This question is asked to check how confident you are in handling your day-to-day work and also your confidence in sticky situations.

You can briefly explain some of the challenges that you faced in your earlier jobs, how you dealt with them, how your team or managers helped you, and how you successfully overcame the problems finally. Avoid talking about bad experiences. Also avoid blaming anyone or the team for any problem. It is better to talk about technical challenges and problems.

22

Are you a team player?

Sometimes the interviewer asks you this question directly. This question is also asked indirectly, as in: 'Do you prefer working in teams or working alone?' or 'How comfortable are you in working as a member in a large team?' This question is to check how good and comfortable you are in working as a team player (particularly in a few types of jobs where team work is very important).

Obviously we need to say, 'Yes', but support your answer with more details or by giving a few instances in the past where you worked very well as a team player. If you are a fresher, you can talk about your participation in or organizing of team sports events, get-togethers, etc.

Focus more about team strength than about individual abilities. This question could also lead to questions like how you handled conflicts within the team. So be prepared!

19

What do you think of your previous boss?

This question is to check how well you can work with, or relate to your new boss if you get the job.

Speak about a few good things you found while working with your previous boss. However, you can't be too open in answering this question!

How much of a salary hike are you looking for?

23

Obviously, this is one of the most difficult questions to answer!

If you're honest and say, 'Double the current salary', you won't get the job. If you say, 'I am fine getting even the old salary', you might actually end up getting it! A safe answer is, 'The same as the industry average hike one gets while moving to a new job' (whatever that 'industry average' means!). If you've done enough analysis about the salary structure in the new company and know that you'll get more for the same level of experience and skill set, you can say: 'The same as the salary that a person with similar experience and skills will get in your company,' and throw the ball back in the interviewer's court.

Why should we hire you?

24

This is a question that every interviewer has, while interviewing a candidate. They want a justification for why they should select you. The interviewer just bounces this ball to you and checks how you give the reason for hiring you!

Tell them about your professional and personal strengths, relevant job experience, or academic background, your suitability for the current job requirements, etc., and give your view on why they should hire you. Bad answers: 'Because I am desperate for a job'; 'I have searched for jobs for more than a year and I didn't get any—you should help me!'

By: S G Ganesh

The author is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill. You can reach him at srganesh@gmail.com.


This article is an excerpt from the book *Cracking the C, C++ and Java Interview* by S G Ganesh. **Courtesy:** Tata McGraw-Hill Education Private Limited, 7, West Patel Nagar, New Delhi 110 008

Do you have any questions for us?

25

Typically, an interviewer will ask this question just before the end of the interview. This is to check if you have any important questions that you want to get clarified. Instead of saying 'I have no questions', it is better to ask relevant questions to show your keen interest in getting the job.

Do show enthusiasm about the new job and ask about the new team, opportunities, company, etc. Good examples: 'What are the current problems that the team is facing now and how can I possibly help?', 'What are the career growth opportunities available in the company'. Bad examples: 'Did I do the interview well?', 'Will I get this job?' (Both are in the list of Frequently Asked Wrong Questions—never ask these questions in the interview! But yes, you can ask 'When can I expect to hear from you?')



Cracking the C, C++, and Java Interview

Author	:	S G Ganesh
Publisher	:	Tata McGraw Hill Education Pvt Ltd
ISBN-13	:	978-0-07-007792-8
ISBN-10	:	0-07-007792-4
Pages	:	264
Price	:	Rs 250/-

TechnoMail - Enterprise Email Server
Anti SPAM, Anti Virus, Email Content Filtering

Firewall, Internet Access Control
Content Filtering, Site Blocking

Bandwidth Management System

Managed Email Hosting Solutions

TechnoInfotech®

1, Vikas Permisses, 11 Bank Street, Fort, Mumbai, India - 400 001. Tel.: 91-22-6633 8900 Ext. 324. info@technoinfotech.com

slackware 13

Is it Really Worthy of the 21st Century?

The latest version of the oldest surviving GNU/Linux distro comes with updates galore, including a KDE4 desktop and ext4 filesystem. It perhaps even has everything for a geek. However, whether it makes the cut for a typical desktop user is the real question.

*S*lackware is old. No really, it's OLD! If you're used to the likes of Ubuntu or any of the other modern GNU/Linux systems, setting up Slackware would seem like you're back in the early 90s.

Some will argue that Ubuntu and the rest that try to make life easier for desktop users don't give *that* fine-grained control over the way you set up your system. Why should a distro try to make choices on your behalf? By configuring your system from the bottom up, you'd ultimately learn the internals.

Does it help? Yes, it does. Immensely! Does everyone have the time and inclination? NO!

Whatever said and done, the majority of those who use computers want to get on

with their work—most don't have the time to peek into the innards of an OS. So, if you belong to this category, you're probably better off without Slackware.

However, Slackware 13 has one surprise in its bag—the desktop somehow looks pretty modern, and works out-of-the-box. Yes, KDE4 has finally infected the conservative Slackware developer(s). Pat Volkerding writes in the release notes, "I'm using it on all my own machines (including an Intel Atom with compositing enabled), and I've really fallen in love with it once I got used to it. The tools are integrated better with the desktop, Qt4 seems to be a faster and more stable platform, and nearly everything that was available for KDE3 has been ported to KDE4 and works great."

Now, Pat doesn't put anything inside Slackware unless and until he's completely satisfied with the piece of code. He'll still make you use LiLo instead of Grub, by default! So, all this praise coming from him could only mean how impressive KDE4 is and what lies ahead in the coming years.

Installation and configuration

I think the first Slackware version I had tried was 9.1—that was in 2003. As far as I can remember, things pretty much look the same, even today. It's the same old ncurses-based wizard that's fool-proof and stable.

Anyway, for the uninitiated, I'll try to walk you through the Slackware path. After booting the CD, you're prompted on whether you wish to use a US English keymap (for your keyboard) or not. We Indians do, so simply hit *Enter*. This brings you to the login screen.

Log in as the root, and type *setup* to start the installation. Note that if you need to configure your partitions, the installer won't offer you anything. So, do your partitioning tid-bits using the *cfdisk* utility and then proceed with *setup*.

This will bring you to an ncurses-based menu. Read what the *Help* option offers if you like, or jump straight to the *addswap* option. Follow the instructions from there on, and you should be up and running.

A few things I noticed:

- Now that the kernel version is upped to 2.6.29, Slackware offers ext4 as the default filesystem. Of course, other choices like ext3, XFS, JFS, etc, are still available.
- It's better if you choose the recommended package installation sections. Although this will install everything available in the DVD, it won't bug you with all those prompts.
- The package installation on my C2D 1.6GHz HP 550 laptop with 1GB of RAM took only 15 minutes. Considering that the total install size is more than 4 GB, that's pretty fast.

General configuration and desktop

Slackware still boots you to init level 3—that is the command line interface with networking enabled. You need to start X with either *startx* or by launching the display manager—for example, *kdm*. This setting is nice if I were to run a server on it, but I'd rather have the desktop show up by default. So, it's time to open */etc/inittab* and set the default runlevel to 4—yes, it's not 5, like other distros.

You'll also need to set up a normal user account, manually. Unfortunately, that was not a part of the installer's job. Anyway, run the following command to create one:

India's Premier Linux & OSS Academy is Offering Online Training. Register Now .. Save upto 25%

Zend

**ZCE - Zend Certified Engineer
(International Cert for PHP)
Zend Framework**

RHCE

**RHCT - Red Hat Certified Technician
RHCE - Red Hat Certified Engineer
RHCS - RH Certified Security
Specialist**

LPI

**LPI 101 Certification
LPI 201 Certification
Ubuntu Certification**

Other Courses :
Shell Scripting
Perl Scripting
Perl - CGI Programming
LAMP Pro (PHP+MySQL)
Python - Django
Ryby on Rails
jQuery
Drupal
Secured PHP Application
Virtualization

Instructor-led online training courses will help you learn OSS quickly and easily. Starts from 14 Oct 09 - CALL FOR DEMO.

**044-42171278
9952966527**

contactus@lynusacademy.com
www.lynusacademy.com



Lynus Academy Pvt Ltd

1961-B, Vijaya Complex - 5th Floor
Asiad Colony, Anna Nagar - West
Chennai 600101 INDIA



useradd -m slacky

Replace *slacky* with whatever user name you prefer.

Logging into the KDE4 desktop made me realise that Slackware has finally bought into the idea of desktop effects—KWin effects work out-of-the-box if your VGA supports 3D. It could also be because Slackware has a strict policy of not customising the defaults offered by upstream software projects and here KDE offers desktop effects by default.

Apart from that, the desktop is plain—with only a panel at the bottom, and no extra icons (nor Folder View) on the desktop. Although Slackware has the Intel Wi-Fi drivers for my wireless Ethernet card, it offers no NetworkManager for seamless wireless access point switching. In fact, you'd need to edit files manually or use command line tools like *ifconfig* and *iwconfig* to set up the network.

For a while, I searched online to find a Slackware NetworkManager package, but unfortunately, none of the third-party Slackware software repos (like linuxpackages.net or slacky.eu) had any software available for version 13 as I write this. Hopefully, things would have changed by the time you install it.

Anyway, the */extras/* section of the DVD has a GTK tool called Wicd, which is a drop-in replacement for NetworkManager. While we're on the subject of GTK tools, they all look ugly due to the absence of the gtk-qt theme engines.

Firing up Wicd gave me a permission-denied error message. The good thing is that the error message also gave a tip on how to fix it. And while we're on the subject of error messages, here're a couple more:

- KMix doesn't load because the user doesn't have rights to the sound device.
- Clicking on an unmounted volume in Dolphin displays error messages.

So, I fired up Konsole to fix things. Adding your user name against the following group names in */etc/group* fixes the aforementioned issues:

- *audio* – this fixes the sound issue.
- *video* – although video worked, I still added my user name here for the heck of it.
- *cdrom* – looked like this solved the CD/DVD mounting issue—though I'm not sure.
- *plugdev* – this was to enable the user to mount other volumes (which probably includes optical discs as well; so maybe adding the user name to the *cdrom* group wasn't required).
- *netdev* – well, this fixed the issue with Wicd permissions.

After saving the file with the changes, and a reboot, Wicd and KMix started as soon as I logged in to KDE, and now I could also mount volumes (hard disk partitions and media devices) using Dolphin. What a relief!

Apps and tools

Finally, coming to productivity and entertainment applications, there're a lot—more than what I needed, at least. And because I went for a full installation, I had more choices for each category of tasks than I'd consider healthy. Anyway, various audio and video file formats worked just fine—even in default KDE apps like JuK and Dragon Player. Yes, there's Amarok and MPlayer (as well as Xine UI) for power users—I missed SMPlayer though. Only thing you'd have to do manually is get the Flash Player from Adobe, Gnash, or elsewhere.

Talk about an office productivity suite and this is where Slackware is a BIG let down. It only offers KOffice version 2, instead of OpenOffice.org 3. I'd like to report here that I never really had issues with the earlier versions of KOffice. But ever since version 2 is out, I find its font rendering capabilities absolutely pathetic—the characters/alphabets look 'fat' with a reddish tinge on black fonts.

Anyway, the good thing is that the accompanying LFY CD has OpenOffice.org 3.1, and you can install it by following this guide: http://wiki.services.openoffice.org/wiki/Documentation/FAQ/Installation/How_do_I_install_OpenOffice.org_on_Slackware%3F


It works like a charm, but the OOo application icons in the menu and run dialogue boxes will be missing. And since UK and US dictionaries don't get installed—you can make use of the OOo Extension Manager (utility located under the *Tools* menu) to find and install them.

Thankfully, the version of Firefox is 3.5, and I also have Ktorrent, Kget and other assorted Internet apps that I require. The GIMP is also available; however, DigiKam is missing.

Well, that basically completes my requirement factors. If you care about programming languages, server-type software, games et al, updated versions of most are available when you go for the default installation. There's even the XFCE desktop for those who like GTKish stuff—sorry, no GNOME here.

As for bugs, the only one I came across was that sometimes, after logging in, pressing Alt+F2 launched the XFCE run dialogue instead of KRunner. A re-login solved this issue. I have no clue why this happens, but it's pretty annoying.


At the end of the day...

I think Slackware scores because of its stability and a lower dependency on memory compared to the rest of the distros. However, it requires plenty of manual configurations before one can be productive. If you can get along fine with that, I guess you won't have much to worry about.  **END**

By: Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau.

Telepathy: Let There be Empathy!



What does Empathy, powered by the Telepathy framework, have in its bag for us? Well, besides audio/video chat, there's geo-location and collaboration, with more coming up.

One of the primary issues faced by new Linux users is the lack of an efficient program for audio/video conferencing along with good integration with other desktop applications. This is something that Mac OSX has done well, with its iChat program. While programs exist for multi-protocol text chat on Linux, until recently we did not have a nice solution for audio/video chat and collaboration. The solution to all these problems requires a powerful communications framework, which is exactly what Telepathy aims to be. Telepathy is as revolutionary a communications application as GStreamer was among media applications.

What's this Telepathy?

The Telepathy project was started in 2005 by a company called Collabora Ltd. It is basically a real-time communications framework and uses the DBus messaging system. It was designed with instant messaging and VoIP-like applications in mind. What makes it really powerful is that it was designed from the start to be highly modular and extensible. It provides a standardised interface, making it easier to implement new features and support new protocols. This also results in code that is easy to maintain and that is portable. Hence, Telepathy can be used on various hardware devices running a variety of software platforms. It is because

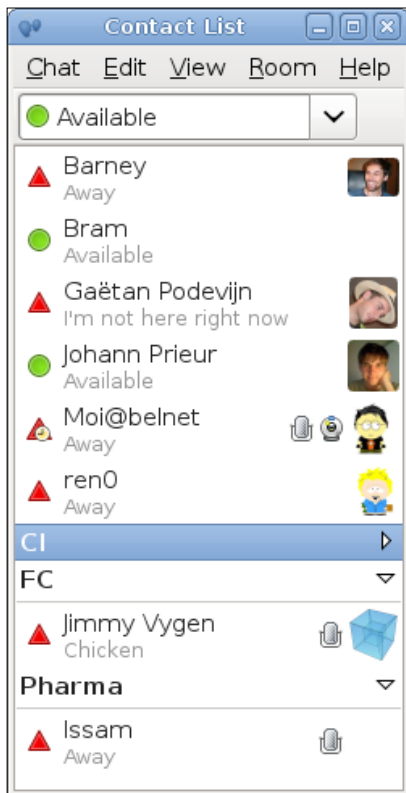


Figure 1: The contact list in Empathy

of these reasons that it is possible for both GNOME and KDE to use the Telepathy framework on their respective desktops with different front-ends.

Enter Empathy!

On GNOME, Telepathy is used by the Empathy IM client. Empathy has a user interface based on Gossip. Empathy supports various protocols like XMPP, Google Talk, MSN, IRC, SIP, etc. Furthermore, Empathy can also use *libpurple* via Telepathy-haze and hence support all the protocols that Pidgin does.

Empathy was included as a part of GNOME from version 2.24. Hence it has been (or is about to be) adopted by some of the major distributions like Ubuntu and Fedora as their default IM client. Its inclusion in these distributions has meant replacing Pidgin, which has caused some controversy, as some people feel it is not mature enough. Though Empathy does require a little more polish, it brings with it a huge number of features that more than make up

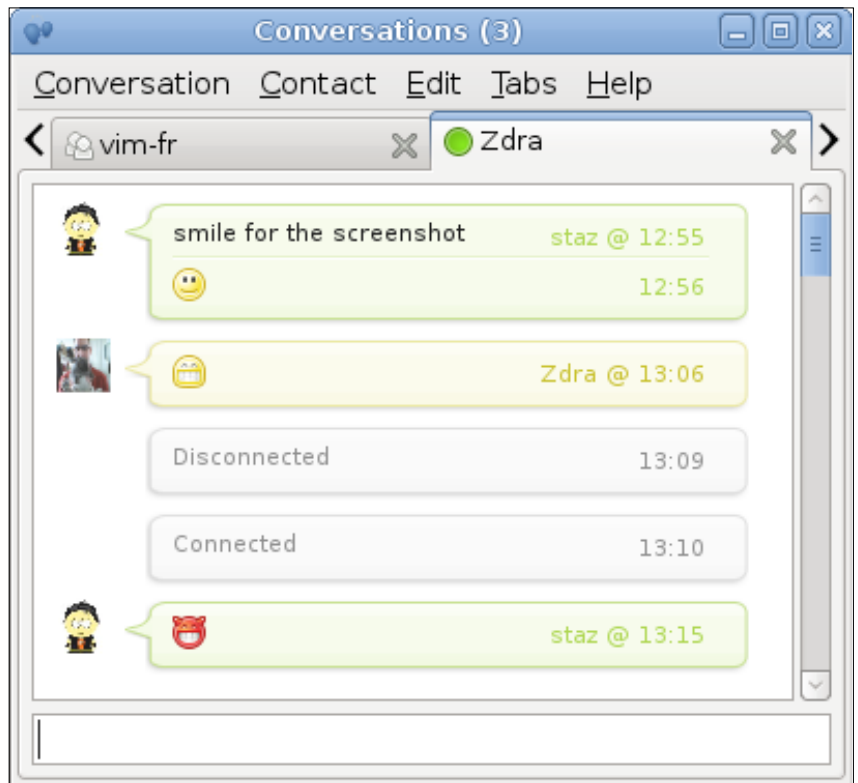


Figure 2: Conversation window with Adium chat themes

for its shortcomings. For example, if you have always drooled over how beautiful Adium looks, then you are in luck as Empathy supports Adium chat themes (and other themes as well), which really make the conversations look great.

Empathy was designed right from the beginning to have good VoIP/video chat features. Currently, video chat works only on XMPP and SIP, though hopefully, MSN and Google video chat should be supported for GNOME 2.28, in time.

Empathy supports video chat using open codecs like Theora and, therefore, should enable out-of-the-box video chat support, which could help us in ditching proprietary apps like Skype. Being able to have audio/video chats over popular protocols will certainly enable a smoother transition to Linux for many people and also improve the accessibility of the desktop.

Another cool feature that Empathy has is the sharing of geo-location information using the GeoClue project. This enables us to announce our current location and

we can, in turn, view the location of all our friends on a map. This is part of a bigger push on the Linux desktop to have geo-location support in as many apps as possible. Location information can range from something as specific as your street address to something much less accurate. The location is determined by using GPS, network information and other available sources that can be used by GeoClue. Unfortunately, this feature is only supported by XMPP servers that support PEP, and sadly, Google Talk doesn't.

A recently added feature to Empathy is the ability to share our desktop with our contacts. We have seen how elegantly this feature is implemented in iChat. A similar simplicity is employed here as well—you just select a contact with desktop sharing capability and that's about it! This is actually implemented using Telepathy's tubes feature that helps set up the connection and uses Vinagre for desktop sharing. This feature opens up a world of interesting applications. Imagine how useful this could be when you have to

fix your parents' computer remotely or want to view files on your home PC from your office.

As Empathy is a part of GNOME, it integrates well with other GNOME applications. For instance, it interacts with NetworkManager to know if a connection is available or not, and if not then it doesn't try to keep on reconnecting. When somebody calls you, it pauses the music in Rhythmbox and then the music resumes once the call ends. It can set your status message to show the music you're currently playing in Rhythmbox, as well.

While all this is certainly very cool and exciting, the best feature that Telepathy brings to the table is something known as tubes. Essentially, what this does is to allow any application to take advantage of the Telepathy framework to add collaboration features with maximum ease. The kind of power and usefulness this feature adds to the applications is tremendous.

We already saw how this will enable desktop sharing using Vino and Vinagre. There are many more applications for this—for instance, we can collaboratively work on a document in AbiWord using the AbiCollab plug-in. We can share currently playing music and other music files, which is soon going to be possible in Banshee (and also Rhythmbox) as a result of a GSOC project. We can also play games with our contacts, a feature that is being worked upon for GNOME Sudoku.

GNOME isn't the only desktop looking to use Telepathy; work is going on so that our KDE brethren can also enjoy this wonderful tool. Kopete is looking to modify itself to be able to use Telepathy's QT port. As Telepathy is a highly modular and portable framework, it is used in many other platforms apart from the Linux desktop. In fact, Telepathy was first commercially deployed on the Nokia 770 to handle its IM and communication needs. Since then, it has been used in many other Nokia devices like Nokia 800, 810 and the


recently released Nokia N900, as it is part of the Maemo platform. Even Intel uses Telepathy for its Moblin platform. The Sugar desktop, which was formerly used by the OLPC project, has also used the Telepathy framework.

Moreover...

Until now, both Telepathy and Empathy have mainly been worked upon to get the basic features in order so that many cool features can be added in the future. Having audio/video chat and file transfer support on all major protocols would really be great. Apart from this, the addition of support for meta contacts using the People or Solynt project is also planned. This will allow us to merge multiple contacts belonging to the same person, on different protocols. Then our contacts list will truly become a friends list.

Better integration with the desktop is being worked upon for Empathy, especially for the upcoming GNOME 3 release. The GNOME Shell is being designed while keeping Empathy in mind. Zeitgeist, the second major application to feature in GNOME 3, will also feature integration with Telepathy to enable easier access to chat history.

To enable audio/video conferences with several people at once, the MUJI (Multi User Jingle) project has been started. This will make it possible to have conferences over an open protocol using open codecs out-of-the-box, from your favourite distribution. All this is so that there is no longer any need for proprietary applications, which seem to be out of place on a free desktop anyway.

I guess with Telepathy, we might have finally solved the communications part of the 'desktop Linux' problem. 

By: Praveen Thirukonda

The author is a third year computer engineering student. He likes playing badminton when not messing around on his computer.



redhat®

TRAINING PARTNER

RHCE / RHCSS Exam Centre

Hundreds of hardcore
Industry Professionals
Vouch for it.



RHCE only @ ADVANTAGE PRO

- Pioneer in LINUX training since 1999.
- Chennai's #1 RHCE training Centre
- 100% results every month.
- Many centum students.
- Also get trained on My SQL - CMDDBA, CMDEV, PHP, Perl, Ruby, Python, Ajax...

Next RHCE/RHCSS Exam
12th & 26th October'09

**Don't wait.
Register Now.**

Let the success vouch for you!



**VECTRA
TECHNOSOFT**

ADVANTAGE PRO

IT TRAINING DIVISION OF
Vectra Technosoft Pvt. Ltd.
Open Source Academy

**Regd. Off: Wing 1 & 2, IV Floor,
Jhaver Plaza, 1A, N.H. Road,
Nungambakkam, Chennai - 34.**
Ph : 28263529 / 3530 / 3540
Telefax : 28263527
E-mail : enquiry@vectratech.in
www.vectratech.in

ODF Olympiad 2009

Winners to Get Netbooks!

It's heartening to hear of the West Bengal state government joining hands with the government of Malaysia and a host of other OSS stakeholders, to drive home the advantages of the ODF amongst students.

The ODF Olympiad (www.odfolympiad.org) was launched in 2006 as a global school-level competition. It started off, on the advice of India's earlier President, Shri A P J Abdul Kalam, with the goal of increasing awareness about open source software in schools. The 2008 edition of the ODF Olympiad witnessed participation from over 116 countries. The competition aims to showcase the ease of using cutting-edge ODF compliant software such as OpenOffice.org and to bring students and teachers closer to adopting free and open source software.

For the 2009 edition of the Olympiad, students of all the participating institutes are required to submit a 21-slide presentation via e-mail, in an Open Document Format (ODF), based on specified topics. And the submissions will be evaluated by an eminent jury, on the basis of pre-determined parameters. The contest will help the government and the schools in faster adoption of the ODF, which would further help in narrowing the digital divide.

What's in store for the winners?

While everyone is assured of a participation certificate, there are prizes for three winners in each of the categories. The winning school will be awarded the ODF Olympiad trophy and there's a Software Freedom trophy for the

school that wins the inter-country contest. What's more, there are also netbooks for the winners, awarded by IOTA (Institute for Open Technology and Application).

The complete evaluation process of the Olympiad has been visually outlined in the flowchart.

Topics for the participants

Category	Standard / Class	Presentation topic
I	Upto 5 th	The benefits of computers to my school
II	6 th , 7 th & 8 th	The benefits of Open Source Software to my school
III	9 th & 10 th	Open Standards: Freedom from the Digital Divide
IV	11 th & 12 th	The importance of Open Standards in e-Governance

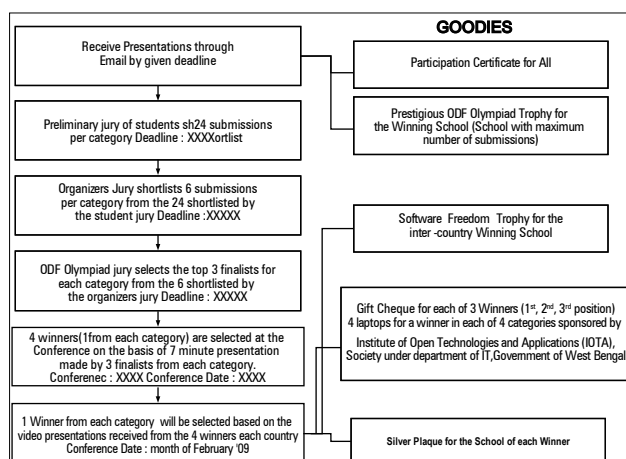
The ODF Olympiad is powered by...

The ODF Olympiad is a multi-stakeholder initiative being promoted by the government of Malaysia, the government of West Bengal, Sun Microsystems, GNU, the Free Software Foundation, the ODF Alliance, IIT Delhi, JNU, IIM Ahmedabad, the Manufacturers' Association for Information Technology (MAIT), IOTA (a society under the Department of IT, government of West Bengal), the ODF Alliance, OSDD, the Knowledge Commons, etc.

OpenOffice.org on LFY CD for Olympiad participants

OpenOffice.org is an office application suite available for a number of different computer operating systems. It is distributed as free software and written using the GTK toolkit. It supports the ISO/IEC standard Open Document Format (ODF) for data interchange as its default file format, as well as Microsoft Office formats, among others. The source code of the suite was released in July 2000 with the aim of reducing the dominant market share of Microsoft Office by providing a free and open alternative.

You can install OpenOffice.org, included with this issue's CD, from <media drive>/software/office/.




Flowchart: Evaluation process of ODF Olympiad 2009





WHO IS GOING TO WIN THEM THIS YEAR?



Supported by:



Categories for EFY Awards*

Industrial Electronics

Resistors	Capacitors	Diodes & Transistors	Relays & Switches	Educational Kits & Trainers
EDA Tools	Solar Cells	Sensors & Actuators	Soldering & Desoldering Stations	Multimeters
Oscilloscopes	Security Systems	SMF Batteries	Inverters & UPS (Industrial)	

Consumer Electronics

Inverters & UPS (Domestic)	EPABX & Key Telephones	DVD Players	Television Sets - CRT	Television Sets - Flat Panel (LCD & Plasma)
Home Music Systems	Mobile Phones	Desktop Computers	Laptops / Notebooks	Microwave Ovens
Water Purifiers	Car Music Systems			

Jury Awards

Electronics Man of the Year	Electronics Company of the Year	Life Time Achievement	Green Company of the Year
-----------------------------	---------------------------------	-----------------------	---------------------------

Sponsorship & Branding opportunities available: e-mail us: efyawards@efyindia.com

16 December, 2009 | India Habitat Center | New Delhi

For more info: www.efyawards.com

GIMP for Beginners, Part 4

The Selection Wand

Welcome to the fourth instalment of tutorials on the GIMP. This is where we start learning about the many techniques and exquisite effects that can be created using this application. In this tutorial, we focus on the selection and gradient tools.

After all the banter about the various tools available, it's time to get our hands dirty with the GIMP. Last month, we discovered various aspects of image editing, along with the interface and layers. But from now on, we will explore new effects, cool techniques and other ways to create exquisite effects. In this tutorial, we will start by creating a very simple wallpaper using just the primary selection and gradient tool.

Before we get started, it is important to know the size of the image we want to create/edit. So if you wish to start creating an image, I would recommend choosing your monitor resolution. Using the same resolution will produce the best effects and you will not need to resize the image, in case you wish to use it as your wallpaper.

Step 1 Fire up the GIMP and click on *File→New*. Set the resolution to that of your desktop. I will be using a resolution of 1920 x 1200 so that those with very high-end monitors can use the wallpaper without any distortion.

Once you have created a new image, notice the white worksheet in the Gimp workspace. (*Ed's note:* You can select the colour of the background and can also make it transparent, if you want.) This is the base for the wallpaper. (See Figure 1.)

Step 2 Now, click on the 'Create new layer' icon to create a transparent layer. Change the layer's name by double-clicking on it or by right-clicking→*Edit Layer Attributes*. Set the new name as *Base*.






Figure 1: Image template

By the way, do not forget to enable a high colour channel, i.e., GEGL. To enable GEGL, just navigate to *Colours*→*GEGL*.

For this tutorial, I will be working on a background with a green blend that has a mix of two different tones. However, you can choose any colour.

Step 3 Set the foreground (FG) colour to #496623 and the background (BG) to #96c756. To set the colour, click on the colour swatches and enter the code in the HTML notation bar.

Step 4 Grab the *Blend* tool from the toolbox or press *L*. Now, with the *Base* layer selected, provide a blend upside down. To achieve this, hold down the left button on the mouse and drag a line from the top to the bottom of the white worksheet, as shown in Figure 3 – part1.

Step 5 Now let's add lines to the wallpaper. Create a new transparent layer and rename it to *Lines*. Now grab the '*Bucket fill*' tool and paint the complete foreground white. With the *Lines* layer selected, head to *Filters*→*Distorts*→*Erase Every Other Row*. Keep the default settings and click on *OK*. Set the layer mode to *Overlay*. See Figure 3.

Once we've got the background done, it's time to add some effects. We will add the effects in three parts of the image, i.e., the top, the bottom left and the bottom right.

Step 6 First, let's create a new transparent layer and rename it *Top Line/Bar*. Grab the rectangle selection tool and create a rectangle as shown in Figure 4 – Part 1. Using the bucket tool, fill it with white. With the same layer selected, create another rectangle just below it but with smaller dimensions as shown in Figure 4 - part 2 and 3. Continue doing the same until you have created three rectangles of varied dimensions. Set the opacity of the *Top Line/Bar* layer to 50. This

gives a smooth faint look.

Step 7 We will now focus on the other selection tool, i.e., the *Ellipse* selection. As usual, start by creating a new transparent layer and name it '*Lower Left Big Circle*'. Grab the *Ellipse Selection* from the toolbox or press *E*. Draw a circular shape at the bottom left portion by pressing *Alt+Shift* and dragging while holding down the left mouse button as shown in Figure 5 -part1. The *Alt+Shift* key combination will help you draw a precise circular shape, without any distortion. Once you are done, navigate to *Edit*→*Stroke Selection*, and with black as the foreground colour, stroke the selection with the line width radius as 1. Deselect by pressing *Shift + Ctrl + A*

Now create another circle with a bigger radius and the same centre (Figure 5 – Part 2). Repeat the stroke effect for the second circle as well. Once that's done, deselect everything by pressing *Shift + Ctrl + A*. Now right-click on the *Lower Left Big Circle* layer

Figure 2: Gradient settings




Figure 3: Setting the initial image



Figure 4: Top border effect

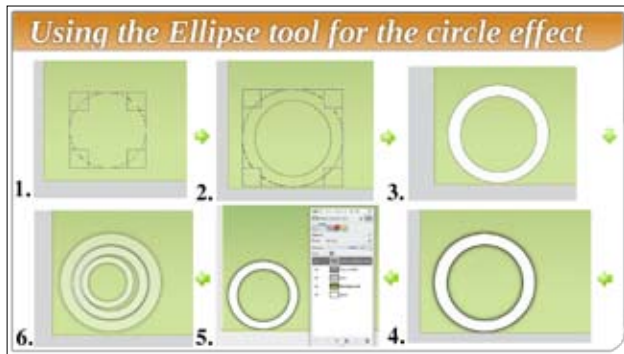


Figure 5: Circle effect

and choose the *Alpha to Selection* option. Notice that the concentric circles become actively selected. The next step is to fill them with colour. We need to fill colour in between their boundaries. To do so, navigate to *Select→Invert*. Fill the selection with white colour using the *Bucket* fill or simply drag the white swatch over the selection.

Right now the circle looks a bit plain, so we will add a shadow to it. The GIMP has a pre-built 'drop shadow' plug-in. Just right click on the Lower Left Big Circle layer, select *Alpha to Selection* and navigate to *Filters→Light and Shadow→Drop Shadow*. A new dialogue box will pop up. Set the X and Y values to 0 and leave the blur radius intact. Once you have provided the drop shadow, you will notice a few transparent parts added. To remove them, just right-click the *Background* layer, select *Alpha to Selection* and crop it using *Image→Crop to Selection*.

The next step is to merge the drop shadow layer. To achieve this, just drag the drop shadow layer above the Lower Left Big Circle layer and select the *'Merge down'* option by right-clicking on the drop shadow layer. Once done, lower the opacity to 40.

Step 8 To achieve a bit of a cartoon-like effect, we will make concentric circles. We need not repeat Step 7 every time. You do remember that we are working on one of the most powerful image editors, right? So, just duplicate the Lower Left Big Circle layer and rename it Lower Left Mid Circle. Resize the layer using Scale Tool from the toolbox (or press *Shift+T*) and make a smaller circle. Use the move tool (m) to move the mid-circle inside the big circle. Continue repeating the step until you have achieved the effect as shown in Figure 5 - Part 6.

I first thought of introducing only the primordial selection tool but later thought I'd add the path tool as well. Don't worry—the path tool work in this tutorial is just to get you familiar with it. We will work with the path tool more extensively in the upcoming tutorials.

The next step will guide you in spicing up the bottom right portion of the image.




Figure 6: Right side effect

Step 9 So, let's get on with the path tool. Create yet another layer and label it 'Right Line'. Grab the path tool from the toolbox or press *B*. Once done, make an inclined shape as shown in Figure 6 (part1). Just keep making dots and complete the selection. To join the two dots, right at the end, hold *Ctrl* and click the nearby dot to complete the selection. Press *Enter* to make the selection active. You will be fretting that I have taken the path outside the canvas. Don't worry, the path tool is very intelligent and will only create a selection inside the canvas. So if you have very little space to work with inside, don't hesitate to exploit the outer territory.

Once you have completed the selection, fill it with white colour and continue making smaller siblings as shown in Figure 6 - Part 2. Fill them as well and reduce the opacity to 60. You can always vary the opacity and the ways to create your own effect, but for this tutorial we will work in a proper manner.

The wallpaper looks a bit simple, doesn't it? So how do we spice it up? Let's add some bubble-like circles to it. We will try to keep it simple with no jazzy effects, as of now.

Step 10 So we are going to add some bubbles. We will work in multiple bubble layers to ease the work. Create a new transparent layer and rename it *Bubbles*. Grab the *Ellipse selection* tool and create a semicircle at the horizon of the wallpaper as shown in Figure 6 - Part 3. Fill it with white colour. Now create a bigger selection, as shown in Figure 6 - Part 4, and navigate to *Edit→Stroke Selection* and stroke the outer part with the line width value set to 6. Deselect by *Select→None*.

The final part is to create another ellipse just above the outer ellipse's border. Note that we have to create a selection just above the previous one so that they look very close and glued together. The newer line should not overlap the previous white line.

Once you have set the foreground colour to #496623, stroke the newly created selection with a line width of 6. This completes the bubble part. You




Figure 7: The final image

must have noticed that the instant you stroke the ellipse, a line appears at the bottom. Also note that the line overlaps the inner white core, the middle white line and the transparent area between them. This doesn't look clean, so we will get rid of it. This is a tricky part and needs some understanding.

Let's start with the inner core. We have been working on the same bubble layer; so if you try to remove the line by using an eraser, you will wipe away the drawing as well. To remove the lining in the drawn part, we will use the brush. Grab the colour picker tool and make sure the bubbles layer is selected. Now click on the inner core. Because we coloured it white, it's not necessary to use the colour pick. But this trick comes handy with unknown colour codes. So grab the brush, whose colour is now set to white, and carefully clean only the area that's inside the inner core. Zoom in for maximum control.

Now we will clean the transparent/background part. Since we are working with layers and we haven't drawn anything in between the inner and the outer lines, we can use the eraser. Press **Shift+E** to evoke the eraser and carefully erase the line from the transparent/background portion, and not from the white lines.

We are almost done now. Just get the brush again and set the FG colour to white. Zoom in and remove the green line from the middle white-bordered ellipse. Continue doing this until you are satisfied. Once done, reduce the opacity to 60. The final look should resemble Figure 6 – Part 7.

This sums up the first bubble/half bubble. Thankfully, we won't have to go through the same procedure again and again. I'll make the next bubble creation simple, since we will be creating the next round of bubbles in the complete canvas.

Step 11 Create a new transparent layer and name it *Main Bubble*. Create a medium circular selection using the ellipse tool and fill it with white (Figure 7 – Part 1). The next step is pretty similar to that of the previous one. We will not go into creating many selections. So to ease your work, we will expand


the current selection. In case your selection went inactive, you can always get it working by right-clicking on the particular layer and selecting '*Alpha to Selection*'. When your selection becomes active, go to *Select→Grow* and increase the selection by 7. Now stroke it with white using the *Edit→Stroke* selection and with 6 as the line width. For the outermost part, increase the selection by 4 and stroke it using 4 as the line width with the FG colour set to #496623. Set the opacity to 60.

To make multiple bubbles, just duplicate the layer, reduce the layer size using the *Scale* tool (**Shift+T**) according to the resolution and move it using the *Move* tool (**m**). Continue duplicating and moving until you reach the effect as shown in Figure 7 – Part 2.

Now, grab the brush and set the FG colour to white. Then create a new layer and name it *Dots*. Varying the brush size from the *Tools* option, create multiple small and medium dots around the circle. You can even use the *Sparks* brush to give a yellow star-shaped light.

Step 12 Congratulations! We are almost done with our first assignment. Now we just need to jazz up the image a bit. Let us quickly create a new layer and name it *Middle Line*. Grab the rectangle selection tool and create a selection as shown in Figure 7 – Part 3. Now choose the *Edit→Stroke* selection and with the line radius as 2, stroke it. Right-click the layer and click on '*Alpha to Selection*'. Navigate to *Edit→Border* and stroke it again using *Stroke Selection* (keep 2 as the line width).

Now grab the path tool and create a path as shown in Figure 7 – Part 4. Hold the vertical line with the left mouse button pressed down and drag it to give it a curved shape. You will notice directional dotted lines, which help you provide a curve to the path. Arrange the dotted line as done in Figure 7 – Part 5. Stroke the path with 4 as the line width. Now create an arrow using either the brush tool or the path tool.

Add some text if you wish, and we are done. Thus, we created a simple, yet funky, wallpaper using nothing but four tools. Isn't it great? Keep trying out new things and do send me the results that you get. In the next tutorial, we will learn how to isolate images from various backgrounds using different techniques. I also have something very special for writers. And do not forget to get the source of this wallpaper from the LFY CD. I have also included some more sample wallpapers. Have fun, GIMPing! **END!** 

By: Shashwat Pant

The author is a FOSS/hardware enthusiast who likes to review software and tweak his hardware for optimum performance. He is very interested in Python/Qt programming and fond of benchmarking the latest Linux distributions and software.

What's Coming in Karmic Koala?

Ubuntu 9.10 is scheduled for release on October 29, 2009. Here's an introduction to the primary features being planned for the release.

May 2009 saw the Ubuntu developers meeting in Barcelona for the UDS (Ubuntu Developers' Summit), where they drew up plans for the next version of Ubuntu, codenamed Karmic Koala and set to be released on October 29. As it is expected that Ubuntu 10.04 (i.e., the version after Karmic Koala) is going to be a Long Term Support (LTS) release, a lot


of changes are planned for this cycle. Which makes this one of the most exciting releases in a long time. Various changes have been proposed at the system level while trying to focus on improving the user experience. Many changes that were postponed in the previous cycles have been accepted this time in the hope of having a better LTS.

Foundation-level changes

Various changes to the core of the system are going to be made, such as:

- **Making *ext4* the default filesystem:** This will lead to a substantial boost in performance as well as reduce *fsck* occurrences by about 10 times. This was present as an option for Ubuntu 9.04 but now will be the default choice for new installations.
- **Moving to GRUB2.**
- **Moving to GDM:** This should lead to much better and prettier login screens. It was not included in previous releases as it lacked configuration options, which are being worked upon in this cycle.

As many changes are being made to the core, I would recommend everyone to reinstall rather than upgrade from a previous release. The reason is that many of these changes will be held back if you upgrade, as it is difficult to update such software without breaking the system. So to get the maximum performance, reinstall from scratch.



Boot experience improvements

It is rightly said that the first impression lasts. Keeping this in mind, Ubuntu is aiming to completely revamp the boot experience. The target for the boot time is 10 seconds on a Dell Mini 9 for the 10.04 version. This would be reduced to around five seconds, along with SSDs and some customisation.

Apart from boot speed, many steps are being taken to make the whole boot experience much prettier and friendlier. The splash screen is being modified to add some much needed bling to it. The GDM is also being changed to make it more polished with various effects.

To achieve both speed and beauty, Karmic Koala will be the first Ubuntu version to have KMS (kernel-mode switching) enabled. It will only work with Intel and ATI graphic cards. Nvidia cards will not be supported until Ubuntu 10.04. KMS will ensure a flicker-free boot experience as well as fast user switching. Fedora already has KMS enabled from the past few versions. So it will be good to have this in Ubuntu too.

New default applications

The default set of apps included in Ubuntu has not been changed much over the past many versions. So the decision to replace Pidgin with Empathy has been somewhat controversial. Empathy is a part of GNOME and uses the Telepathy framework for protocol support, which makes it highly scalable and reuseable. It brings with it a lot of useful and exciting features like audio/video chat, desktop sharing, etc. Currently, audio/video chat is supported only for SIP and XMPP, while, hopefully, support for Google Talk and MSN should be added in time for Karmic. This change would mean we can finally stop using proprietary apps like Skype. Another exciting feature that Empathy has is geolocalisation support, which helps in sharing and viewing location information. Unfortunately, this only works with XMPP servers but not with Google Talk.

With the goal of having a more social desktop, Gwibber will also be included as a default app. This will help to keep in touch with friends via microblogging sites like Twitter and Identi.ca as well as Facebook, Digg, etc.

Gnome-Bluetooth will now be used to establish Bluetooth connections. This should fix many of the problems that users faced in previous versions, while also bringing in some cool features like A2DP support so that we can have out-of-the-box Bluetooth headset support.

While it was initially proposed to replace Rhythmbox with the Banshee Media Player, this plan was dropped due to its development schedule not matching with Ubuntu's.

Papercuts project

How many times have you heard someone say, "Linux lacks polish and is rough around the edges..."? Well, it seems that the Ubuntu developers too have heard this and to fix it, they've started the Papercuts project. This will

target around 100 usability bugs that are really annoying but not too difficult to fix. This project will add the level of professionalism and completeness found in commercial operating systems. The full list of bugs to be fixed under this project can be found at <https://launchpad.net/hundredpapercuts/karmic>.

Ayatana project


As part of Canonical's efforts to improve the user experience on the Linux desktop, the Ayatana project was started. Ayatana is actually a collection of various projects where design and user experience are given high priority. Currently, Ayatana comprises two projects. The notify-osd project, the new notification system that first made its appearance in Ubuntu 9.04, is being improved in this cycle. The other project is the messaging-indicator project that also was first seen in Ubuntu 9.04. This helps to reduce the clutter in the panel by reducing the number of icons. In this cycle, work is on to make it support many more apps like Empathy, Thunderbird, etc.

GNOME 3 preview

Though GNOME 3 is not scheduled to be released until March next year, a preview release of various GNOME 3 apps like GNOME Shell and Zeitgeist will be made available via the repositories. This will help in getting a lot of exposure for these apps and help in a smoother transition to GNOME 3. People will be able to test the apps and give their feedback to the developers.

Android on Ubuntu


With many people interested in Android, Ubuntu doesn't want to be far behind and is working on integrating Android apps into Ubuntu. This will help netbook users to run Android apps on Ubuntu, natively, which will open up a whole new world of apps for Ubuntu users. Ubuntu also wants to become the best platform for development of Android apps. These steps will hopefully help in increasing the number of companies offering Ubuntu on their machines.

While this just about sums up the main changes to be expected in Ubuntu 9.10, many other improvements can be expected. In the next few months we will see the release of many new operating systems like Snow Leopard and Windows 7, and to compete with them will be a tough job for Ubuntu and desktop Linux, in general. But looking at the features planned and keeping in mind how far we have come in just the past two years, there seems to be a lot of hope for Linux on the desktop. Next year could definitely be the year of Linux on the desktop! 

By: Praveen Thirukonda

The author is a third year computer engineering student. He likes playing badminton when not messing around on his computer.

Liberating Young Scholars



In this and subsequent articles, we will try to explore FOSS-based software for education.

Today, computers are closely linked to education. While some schools use them as teaching aids, the more common uses of computers are for 'computer science', 'computer education', or 'computer literacy'. This is all very well except for a couple of aspects. One, sometimes computer activities are incorporated just for their own sake rather than as an analytical and problem-solving tool that extends the human reach. Two, and this is a subject close to our hearts, computer skills at home and school are imparted almost exclusively using proprietary tools. There are some exceptions, such as in the state of Kerala in southern India.

The purpose of education is purportedly to liberate. Paradoxically, it is computer education that needs to be liberated from closed mindsets.

Mindset Number One: Computer skills are about learning computer languages or tools—how often have you asked people about their skill-sets only to be told, 'Java', 'C++' or 'MS Office'? In case we have forgotten, computer skills are about solving real-world problems using computers and computer techniques, period! Knowing a particular language is just a small fraction of the gamut of possibilities computers offer.

Mindset Number Two: Train school students in all the important proprietary software suites and you have made them 'employable' with the requisite computer skills to fit into the world's largest back office—something that trade bodies want India to become as it comes of age. While some would believe this is a great role for India to play on the world stage, I feel that as a country with such vast potential we should set our sights higher.

We do not have a solution for the melting polar caps nor are we discounting the value of certain proprietary software in education. What we are saying, however, is that FOSS has a stellar role to play in revolutionising education in India in a very cost-effective manner. The liberating ideology and community orientation of FOSS is thrown in as a bonus. If you are the programming type, you can also modify the application to suit your requirements; this being one of the four freedoms that the GNU/GPL (as well as many other software licences) guarantees. It is time we gave FOSS a chance.

In this article we will look at how you can help our young scholars make the most of FOSS at home and at school. In the next part we will talk about how FOSS can help with enabling infrastructure on the education front. Our approach will be descriptive enumeration giving pointers for further exploration rather than a detailed review. That way, the fun of further exploration remains all yours! We are going to stick to GNU/Linux though, with no offence to Open Solaris and BSD fans.

The money is not funny!

If you are still on the fence about adopting FOSS, try this story for a fence ripper. At the time of writing this, the Maharashtra State Government has signed (away?) an MoU with a proprietary software giant for ICT deployment and training in its schools. [You can find pointers to this story in the References section. The Bangalore Student Edition of The Times of India dated August 25, 2009 also carries the story.] According to the report, this has an impact on 68,000 government schools with 300,000 teachers and 8 million

students. Critics mention that just buying the software for schools will cost Rs 238 crores (a crore is 10 million!). The logical culmination of this—with a computer for each student and teacher—is an potential expenditure of more than Rs 2800 crores in Maharashtra alone. And we haven't even started talking hardware.

This is the taxpayer's money. India can ill afford such profligacy when 80 per cent of its population lives on less than Rs 20 a day and 26 per cent on less than Rs 12 a day.

Home infrastructure

Unless you have your office footing the bill for a proprietary licence for your home PCs or you really like to spend money buying licences, you should have Linux running on at least one of your machines at home, if not all. If you have not sampled the Linux desktop lately, you are missing something. Mint, Ubuntu (with its flavours Kubuntu and Edubuntu), and Mandriva, among many others, offer a great desktop experience. The productivity tools (e-mail, browsers, chat, office suites, etc) have matured from their early days and what's more, they are all based on open standards. They also interface seamlessly with proprietary formats that matter. Linux as the platform of choice opens the door for FOSS-based education tools, though many of the education tools run on Wintel platforms, too.

Kindergarten

Junior wants to play Daddy and guess what, he wants his own computer time, too. This is just the right time to introduce some mouse training and maybe the alphabet. The following are some good computer discovery programs.

ChildsPlay (schoolisplay.org) is a program for mouse and keyboard with puzzles, memory games, flash cards and sound.

Gcompris (gcompris.net) is a computer discovery program for 2- to 10-year-olds with over a 100 activities. More seem to be getting added regularly. This program has a long list of testimonials and is also being used for rehabilitation and training of differently-abled people.

PySyCache (pysycache.com) is an amazing theme-able and extensible mouse trainer. Your kids will love its look and feel.

Both KDE (edu.kde.org) and GNOME (live.gnome.org) have their own educational projects. KDE and GNOME are two prominent desktop environments in Linux, among many others.

Klettres helps learn the alphabet and sounds in many languages. It even has Hindi, Kannada and Telugu language packs. Klettres is as helpful for adults learning a new language as it is for kids.

KwordQuiz is a vocabulary building program based on flash cards and multiple choice questions that can be used to administer your own quiz.

Parley is another feature-rich, flash card driven learning program that can be used for vocabulary training or general purpose quiz practice.

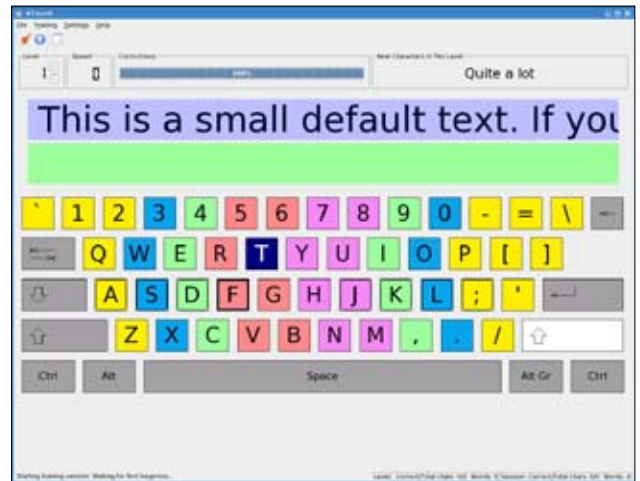


Figure 1: KTouch is a program to learn typing

No Kindergarten software suite is complete without a painting program. *TuxPaint* (tuxpaint.org) is feature packed to keep your tot engaged for hours.

Primary school level

For some advanced vocabulary skills there is *Kanagram*, *KHangMan* and *WordPlay*. Asymptopia (asymptopia.org) has some interesting programs like *TuxWordSmith* and *Crossword Builder* for advanced users.

Now there are math skills to be learned, too.

TuxMathScrabble, *Multiplication Station*, and *Math Crosswords* (all from Asymptopia) can give the young scholar a run around the block. For exercises in fractions there is *Kbruch*.

This is probably a good time to learn typing, too. Given that much of our time (for the gadget freaks, at least) is going to be spent interacting with a keyboard-based device, a hunt-and-peck approach at the keyboard can be a big productivity drainer. *KTouch* and *Klavar* are typing tutors that help you learn and practice at your own pace, with graded lessons.

Logo is a popular language taught in schools to introduce basic programming concepts while controlling a turtle that leaves a trail as it travels across the screen. *KTurtle* implements the Logo programming language with an uncluttered interface.

OLPC

It gets even better. You have probably heard about the One Laptop per Child (OLPC) project. If not, it is worth a visit to *laptop.org* to learn how bright minds are making learning both child-centric and fun using the best FOSS technology on offer today. The OLPC runs a modified version of Fedora, but the attention grabber, or shall we say game changer, is the Sugar desktop.

Sugar is a desktop that articulates a child's world view and has a bunch of activities with learning objectives. They are: word processing, music, calculator, chat, memory games, Logo, and many more applications—all whittled down to a complexity level that is age appropriate. This makes Sugar the ideal one-stop, but with a gentle



Forthcoming Short-term Certificate Courses on

PCB Design, Robotics, Basics of Electronics & Microcontrollers

At Bangalore & New Delhi
EFY Tech Centers

Who can benefit?

Electronics engineering students, hobbyists, working professionals wanting to learn more and job seekers

What do they get out of it?

Job seekers and working professionals add to their qualifications, and students and hobbyists assemble exciting engineering projects on their own



Course duration up to 10 days



First come first served



Limited seats



Don't wait, join now

NEW!

Special weekend course on PIC Micro controllers
from 1 Nov-30 Nov (All Saturdays & Sundays)
Limited Seats, Register Now.

PCB Design (Delhi & Bangalore)

Duration: 5 Days (30 hours); **Schedule:** B'lore: 28 Sept - 3 Oct
Delhi: 12 Oct-16 Oct / 26 Oct-30 Oct / 2 Nov-6 Nov / 16 Nov-20 Nov
Course Fee: Rs 4000/-* **Eligibility:** Engineering students,
working engineers, working diploma holders, hobbyists

COURSE DESCRIPTION:

• Learn to use EAGLE PCB design software by gaining understanding of the basic concepts of PCB design and skills in drawing the schematic, editing and routing and creation of a library • Make a PCB layout of an electronic circuit with various components such as a microcontroller, ICs, transistors, diodes, passive components, switches and connectors

Certificate awarded by EFY

Microcontroller: Microchip PIC16F877A (Delhi & Bangalore)

Duration: 10 days (60 hours); **Schedule:** Delhi: 5 Oct - 16 Oct /
19 Oct - 29 Oct / 2 Nov - 12 Nov B'lore: 9 Nov - 20 Nov
Course Fee: Rs 6000/-* **Eligibility:** Engineering students,
working engineers, working diploma holders

COURSE DESCRIPTION:

• Understand the architecture of PIC16F877A microcontroller • Understand the memory organisation, special-function registers and on-chip peripherals • Learn instruction set and writing of codes in Assembly language • Learn the simulation, real-time debugging and programming processes with the help of MPLAB ICD 2 and PIC Development board • Learn integration of common peripherals with microcontroller through hands-on lab exercises

Certificate awarded by Microchip Inc., USA

Accommodation can be arranged in a private hostel/guest house on prior request.

* **Mode of payment:** Please add 10.3% Service Tax while paying fee. The fee may be paid in cash or through a DD in the name of **EFY Enterprises Pvt Ltd**, payable at Delhi or Bangalore. **You can pay the entire fee at one go or register by paying Rs 1000 as advance towards the course fee.**

Payments should be sent to:

Ms Jayashree, **EFY TECH CENTER**
No. 9, 17 Main First Cross, HAL Second Stage, Indira Nagar,
Bangalore. Phone: 080-25260394, 25260023

EFY
TECHCenter
E-mail: efytech@efyindia.com

Ms Rinku, **EFY TECH CENTER**
D88/5 Okhla Industrial Area, Phase 1, New Delhi 110020
Phone: 011-26810601 to -603

Basics of Electronics (Delhi & Bangalore)

Duration: 5 days (30 hours); **Schedule:** Delhi: 5 Oct - 9 Oct / 26 Oct - 30 Oct / 2 Nov - 6 Nov / 23 Nov - 27 Nov
B'lore: 26 Oct - 30 Oct / 23 Nov - 27 Nov
Course Fee: Rs 4000/-* **Eligibility:** Engineering students, diploma holders, 12th pass with physics, hobbyists

COURSE DESCRIPTION:

• Know about electronic components (both active as well as passive)
• Learn important circuit blocks which will help you in developing your own minor project • Familiarise with basic test equipment for testing and fault-finding • Learn soldering through assembly of an EFY kit

Certificate awarded by EFY

Renesas R8C Microcontroller (Bangalore)

Duration: 10 days (60 hours) **Schedule:** Bangalore: 12 Oct-22 October
Course Fee: Rs 6000/-* **Eligibility:** Done a basic course in C language programming

COURSE DESCRIPTION:

Architecture of R8C microcontroller; C language programming for integration of peripherals; Hands-on training on use of Renesas Starter Kit, board, E8 debugging emulator and application board in HEW environment(HEW integrates various tools such as compiler, assembler, debugger and editor into a common GUI). Learn to evaluate Renesas microcontroller performance for your product application


Certificate awarded by EFY

Robotics (Delhi)

Duration: 10 days (60 hours); **Schedule:** Delhi: 9 Nov- 20 Nov
Course Fee: Rs 8000/-

COURSE DESCRIPTION:

Science and engineering of robotics; electronic and mechanical hardware related to robotics; basics of PIC16F877A applications used in robotics. Project work: Construction of a robot based on PIC16F877A MCU TAKE HOME YOUR OWN ROBOT AFTER COMPLETION OF THE COURSE



DRIVE OR BE DRIVEN



THE ONLY
IT SUPER HIGHWAY
IS NOW OPEN

November 11-13, 2009



CHOOSE YOUR CONVOY

This edition, we give you three packed days of knowledge and opportunities. The focal theme for this year is "ICT for a Billion People". So choose from nine special tracks focusing on various industry verticals, a trade show with five focused pavilions, session on IT for Green & E-waste and Good People Practice, a CEO-CIO Conclave, a special Awards Function and various social and business networking platforms.

Multitrack Conference

International Exhibition

CEO/CIO Conclave

Awards Function

Bangalore IT
Power Golf Tourney

Workshops

International
Conferences

CHOOSE YOUR VEHICLE



The plenary talks will deliberate on the environmental concerns to be addressed through IT for Green & E-waste management and social issues of GPP (Good People Practices) and Empowering Women in IT. The IT Vision 2020 Leadership Series will enlighten young minds about the Emerging Opportunities in the ICT industry.

The Conference will be addressed by experts from ICT driving technology & business as well as policy makers from the Centre and leading States of India.

So, choose the vehicle which will you give you the best mileage and make the most of your drive!

1

IT facilitates
GOVERNANCE

2

IT powers
ENERGY

3

IT navigates
GIS

4

IT cares for
HEALTH

5

IT ExcITE's
AUTO INDUSTRY

6

IT secures
FINANCE

7

YESSS: Opportunities for
10 Startup Companies to
meet Industry & VC's

8

IT supports
EDUCATION

9

IT connects
TELECOM

TURBO FUEL OF THE EVENT

CEO-CIO Conclave: Networking goes a road further, with the CEO's now connecting with the end users of the IT industry – The CIO's. By invitation only.

Awards Function: A trophy for those who cross the chequered flag! The Awards Function is our way of congratulating the best IT performers for the year & Technovation Awards 2009 by ISA.

Rural IT Quiz: A hit from the past, the Rural IT Quiz is the BangaloreIT.biz special, dedicated to the talents all over Karnataka.

IT Power Golf Tourney: Mixing business with pleasure, BangaloreIT.biz Golf Tourney is the express-way for social networking! By invitation only.

Event Secretariat: BangaloreIT.biz

MM Activ

9, UNI Building, Ground Floor, Thimmaiah Road,
Vasanthnagar, Bangalore 560 052, India
Tel +91 80 4092 6211/12/13.
Fax: +91 80 4092 6214

Email: enquiry@bangaloreit.biz
www.bangaloreit.biz

Media Support



Event Partner



Event Manager



You Can Bank On Open Source!

IDBI Bank, with its 506-branch national network, chose to explore the opportunities offered by open source software pretty early. And almost a decade later, the bank feels open source has helped reinforce its image as a customer- and technology-savvy bank.

*B*anks, like other enterprises, have opened their arms to open source-based technology for its obvious cost advantage. But talk about deploying Linux for critical applications and these very financial institutions seem to nurture numerous doubts about security, maintenance and reliability. However, back in 2003, IDBI Bank quelled all such apprehensions by placing critical applications like Oracle ERP, comprising Oracle Financials and HRMS Modules, on a Linux platform. With its roots going deeper into the open source soil, IDBI Bank's success saga remains one to learn from.

The first shoots

In 2001, IDBI Bank planted its first open source sapling by deploying Linux for its enterprise e-mail server and IVRS applications. And what prompted them to shift to open source? IDBI Bank had noticed open source technology fuelling the growth of the Internet with key applications such as the Apache Web server, the Sendmail mail transport agent, and WebStore (open source e-commerce); languages like Java and Perl; and mark-up languages like HTML, WML and XML. The bank wanted to use the potential of open source for the effective management of its processes.

"We decided to evaluate open source as it would help us lower information technology costs, implement flexible software solutions, achieve better security and get software developed using collaborative and cooperative values and not proprietary values, which would be apt for our local and organisational needs," says Sanjay Sharma, MD and CEO, IDBI Intech, a fully-owned subsidiary of IDBI Bank.

IDBI Intech provides IT-related services to the IDBI Group companies and to other organisations, focusing mainly on the BFSI sector. It provides IT services



(end-to-end core banking consultancy), information security services, business process outsourcing, and learning and development services.

One of the immediate results of adopting open source was cost savings. IDBI Bank deployed Linux-based applications without incurring any additional costs other than that of the application. "The applications deployed using Linux were highly cost effective as the configuration of hardware required is low compared to other operating systems. We have also saved on software licence costs by deploying solutions on Linux," says Sharma.

Within a short period of a year or two, open source cut IT expenditure tremendously. "A leading bank in India that adopted open source in more than a thousand branches saved about \$1 million in two years only by way of licence fees, training and maintenance cost," testifies Venkatesh Hariharan, corporate affairs director, Red Hat. Apart from cost savings, scalability, and stability, the availability of a lot more versatile tools was another benefit enjoyed by IDBI Bank.

However, transition hitches are a common anxiety faced by most enterprises adopting open source technology. The case with IDBI Bank would have been no different if the company hadn't decided to deploy applications on Linux, early. Further, it did not completely abandon proprietary software and went in for the change gradually, allowing people to adapt to open source more easily. "Workstations run mainly on the Windows operating system and the use of Mozilla's applications and Star/OpenOffice.org on the Windows platform is quite common," says Sharma.

Open source takes deeper root

Over the years, IDBI Bank has only dug deeper into the open source crust to be in sync with e-age banking and offer customers the latest applications. "To meet the growing business requirements, we needed to deploy new applications. IDBI gave preference to applications ported to open source platforms. Our critical applications have been deployed on open source," says Sharma.

In 2002, the bank launched the open source Jabber application and in February 2003, it became the first bank in the Asian Pacific region to implement ERP

on Linux. Subsequently, many applications have been implemented on open source and today IDBI uses Linux for critical applications related to enterprise e-mail, and Web, DNS and LDAP servers.

The deployment of Linux for mission critical applications at IDBI emphatically proves how successful open source technology is, as there is no room for error in such applications. "The entire electronic communication in our bank is dependent on the e-mail server based on Sendmail, which is on

Linux. If, for any reason, the mail server is down, the communication within the bank and outside will be critically hampered, thus affecting business. This application is very critical for our bank and we have chosen open source to implement it," says Sharma, underlining his confidence in OSS for mission-critical applications.

Apart from mission-critical deployments, open source has helped IDBI beef up security and the confidentiality of information. "Contrary to the common perception, open source has helped keep our systems secure as patches/upgrades are available easily. There are more people inspecting the source code, finding and fixing possible vulnerabilities. Unlike in proprietary software, we don't have to accept the level of security the software vendor is willing to deliver and the corresponding pace at which they release patches and updates," says Sharma.

Yet, a lot of doubt exists about the long-term support and maintenance of an open source IT environment. IDBI Bank took care of this by entering into a support arrangement with Red Hat Professional Consulting for support services like implementation, OS tuning services, consulting and engineering services through six

proactive onsite visits. Further, it has the expertise of IDBI Intech for internal guidance as the bank marches down the open source road.

A growing tree

"Altogether, open source deployment has enhanced productivity and the efficiency of processes and people both directly and indirectly, positively impacting our



"We decided to evaluate open source as it would help us lower information technology costs, implement flexible software solutions, achieve better security and get software developed using collaborative and cooperative values"

—Sanjay Sharma, MD and CEO, IDBI Intech

IDBI's Linux portfolio	
Year	Linux deployment
2001	Enterprise e-mail server and IVRS applications
2002	Jabber for messaging
2003	Oracle ERP on Linux
2009	Web Server, DNS and LDAP server

customers and reinforcing our image as a customer- and technology-savvy bank," vouches Sharma.

Looking ahead, IDBI is in the process of implementing a new version of the e-mail solution, including archival solutions based on open source architecture. "The set-up includes cluster and virtualisation technology on Linux. Users will be able to access e-mail from the Web and smart phones. They can also access the calendar and contacts on smart phones, which can be synced with PIM (personal information manager) applications," says Sharma.


A 'thumbs up' for open source in banks

Thus, it's a 'thumbs up' to open source deployment in banks, especially now with more vendors offering/porting solutions on open source.

Advising other banks/institutions deploying open source for business critical information, Sharma says: "Banks must increase the use of open

source in the front office, where it is still at a nascent stage. Banks can use Eclipse, an integrated development environment (IDE) for Java, to develop complex trading/internet banking systems. ERP applications on open source platforms can cut costs and achieve stability. Further, there are open source data analytics software and a host of financial applications. Linux also supports advanced networking features. Besides a reliable TCP/UDP/IP protocol stack, features like firewalls, quality of service (QoS), tunnelling, etc, have been added to the OS kernel."

QoS support in Linux encompasses a description of the differentiated services effort, firewall implementation using IP chains, VPN implementation using GRE tunnels and advanced routing implementation using Netlink sockets.

Guess it's high time banks stopped hesitating and started banking on open source, not only because of greater reliability and cost savings, but for the benefits it offers its customers. **END** 

By: Vanisha Joseph

The author loves to experiment and writing for LINUX For You is her latest experiment. So, beware! Just a minute, she also happens to be a journalist during the day.



Calling...

Share your knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of linuxforu.com we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at

lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developers or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools



S.G. Ganesh

The 'Struct Hack' Technique

This month we'll discuss an interesting programming technique that started as a 'hack', entered mainstream programming and then became part of the C language!

What strikes you first when you see a structure definition like this?

```
struct array {  
    int nelem;  
    int data[1];  
};
```

If you are not familiar with the technique behind the code, you might think that it is crazy. Why would anyone create an array of size 1, as in "int data[1];" here? We can as well declare it as "int data;," right?

But if you go ahead and check how the "array" struct is used, you can see some still stranger code:

```
#define ARR_SIZE 10  
  
struct array *ptr;  
ptr->nelem = ARR_SIZE;  
ptr = malloc (sizeof (struct array) + sizeof (int) * (ARR_SIZE - 1));  
  
for (int i = 0; i < ptr->nelem; i++)  
    p->data [i] = i;
```

Can you guess what is going on now?

It is not possible to create variable size arrays in C. We can always dynamically allocate an array, but often it is convenient to have a structure where we can have one member as a variable size array.

A 'shrewd' way to overcome this limitation is to have an element in struct to remember the size of the array we want and then have an array with one element size. Then, instead of allocating memory dynamically as "sizeof(struct array)", we calculate the size to include the amount of memory we want to allocate for the array also. But the array already had one element, so we need to adjust for that while dynamically allocating memory (it is not possible to have zero length arrays in C; it is possible in languages like Java).

We have calculated and allocated enough memory for the array and access only within the limits. The dynamically allocated block is contiguous, and compilers usually do not check for the 'out-of-bounds' access, so it


works fine in almost all C implementations!

What we just saw is called the 'struct hack' technique (also known as the 'one-element-array-trick'), which started as a 'hack' to overcome the limitation of the fixed length of structs/arrays. There are some variations; for example, instead of "int data[1];", the declaration can be "int *data;". Since this technique worked well for most implementations and it solved a practical problem in many C applications, it started to get used widely not as a hack but as a way of overcoming the C limitation. However, strictly speaking, accessing past the end of the array results in 'undefined behaviour', according to the original language definition. So, this technique was later made as part of the language.

The C99 standard added a feature called 'flexible array members', which is nothing but the 'struct hack' technique with slight differences. Instead of allocating the last member of the struct as an array of size 1, the standard allows to declare them with no mention of the size. For example, here is the same 'struct array' we saw earlier, written in C99:

```
struct array {  
    int nelem;  
    int data[]; // FLM  
};
```

The dynamic memory allocation can be done using malloc, as described earlier. But it will not consider the flexible array member while calculating the size of the struct, which makes the expression in the malloc easier.

Note that this technique is only for C. For most other languages there is no need to use this technique. For example, since C++ has better facilities, we can use vectors for variable sized (dynamically growing) arrays.  **END**

About the author:

S G Ganesh is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill. You can reach him at sgganesh@gmail.com.



Tips & Tricks



Add extra swap space quickly

Check your used swap space using the following command:

```
free -m
```

or:

```
cat /proc/swaps
```

If the system has very little swap space, you can add more swap easily. Take the following steps:

- 1) Create a file that is of 1 GB size:

```
dd if=/dev/zero of=/dev/myswap bs=1024 count=1024000
```

- 2) Enable this file as swap space:

```
mkswap /dev/myswap  
swapon /dev/myswap
```

This will create an additional swap space of 1 GB. You can see the new swap added as:

```
# swapon -s
```

Filename	Type	Size	Used	Priority
/dev/sda6	partition	1020088	0	-1
/myswap	file	1016	0	-2

Append the following entry to */etc/fstab* to make this swap permanent:

```
/dev/myswap swap swap defaults 0 0
```

—Govindarajalu, govind.rajalu@gmail.com



Terminal magic

Suppose you have to run many applications on a terminal (say, on the GNOME terminal) and you want to automate this process, you can use the command-line options of the GNOME terminal. For example:

```
$ gnome-terminal --geometry=80x20+700+0 --title=ping -x ping google.com &
```

...where:

- *geometry* will decide what the length and width of your terminal will be, as well as its position.
- *title* sets the title of the terminal.
- *-x* is for the command to run.
- *&* to run this as a background process.

See the man pages of the GNOME terminal for more options. You can set tabs and many more options by consulting the man pages.

Now, assume we want to run the following applications, viz.:

- 1) my server—let's say the name is */home/xyz/myServer*
- 2) *tail -f /var/log/messages* to monitor the logs
- 3) *ping google.com*
- 4) SSH to some remote machine—for instance *xxx.yyy.zzz.aaa*

So I will make a script called *automate.sh* with the following information:

```
#cat automate.sh
```

```
gnome-terminal --geometry=80x20+0+0 --title=myServer -x ./home/xyz/  
myServer &  
gnome-terminal --geometry=80x20+0+400 --title=logMessages -x tail -f  
/var/log/messages &  
gnome-terminal --geometry=80x20+700+400 --title=ping -x ping google.  
com &  
gnome-terminal --geometry=80x20+700+0 --title=remoteMachine -x ssh  
root@xx.yy.zz.aa &
```

Now you can run all the applications with a simple command like...

```
$ sh automate.sh
```

Please remember that the opened window will close as soon as the command that is running stops executing. For example, the window/terminal running `tail -f /var/log/messages` will immediately shut off if you don't have permission to read `/var/log/messages`.

—**Shiv Premani**, shivpremani@gmail.com



Back up MBR with the `dd` command

Here's how you can do it:

```
# dd if=/dev/sdX of=/tmp/sda-mbr.bin bs=512 count=1
```

Replace *X* with the actual device name, for example, `/dev/sda`.

And here's how you can restore the partition table to the disk:

```
# dd if= sda-mbr.bin of=/dev/sdX bs=1 count=64 skip=446 seek=446
```

—**Remin Raphael**, remin13@gmail.com



Copy, paste and cut operations in the Vim editor

In order to copy a line of text in Vim, place the mouse on the line and press "Esc yy". If there is more than one line from the mouse pointer then prepend yy with the number (of lines) like this: "Esc 3yy". This will copy three lines. Observe the statement "*n* lines yanked" (where 'yanked' means copied) at the bottom of the file.

Now, to cut a line, place the mouse on the line and press "Esc dd". If you want to cut more than one line, e.g., three lines, press "Esc 3dd".

Finally, to paste, move the mouse to one line above the row where you want to paste, and press "Esc p".

—**Indraveni**, indravenik@cdac.in



Vim tips on auto completion

When cranking out code, or even writing docs, Vim can do a lot for you. It can do word/pattern/line completion using `^x ^n|^p` and `^x ^l`. (Note that the `^x` character means you have to hold down the Ctrl key and then press x.)

For example, open an existing 'C' source file in `[g]vim`, go to insert or append mode, type `pri^x^p`, and you should get a pop-up showing word completion choices which includes `printf` (and `printk` in a kernel module!).

You can also do `pri^x^l`, which will give you an option to select complete lines that start with 'pri' to auto complete.

Important caveats:

1. You need to be in insert mode.

2. The word/line completion appears to work only with 'existing' words (or lines) in the current file. (So if you don't have a `printf` in there, `pri^x^p` won't complete the 'pri' snippet to `printf`).

Nevertheless, it's still very useful!

Resource: <http://www.thegeekstuff.com/2009/01/vi-and-vim-editor-5-awesome-examples-for-automatic-word-completion-using-ctrl-x-magic>

—**Kaiwan Billimoria**, kaiwan@designergraphics.com



CLI calculator

`bc` is an arbitrary precision calculator language.

You can use it by simply typing `bc` in a command prompt and then query your calculation like `1+2` or `3-2`, etc.

You can also use it in your shell scripts. For example, if you want to find the square roots for numbers, then simply create a file called `sqr.sh` and add the following content:

```
#!/bin/bash

if [ $# -ne 1 ]
then
    echo 'Usage: sqrt number'
    exit 1
else
    echo -e "sqrt($1)\nquit\n" | bc -q -i
fi
```

Now make the file executable by using the following:

```
# chmod +x sqr.sh
```

That's all. You can test it now by trying to look up the square root of 36.

```
$/sqr.sh 36
sqrt(36)
6
quit
```

—**fayiz mustahafa**, fayizk1@yahoo.com



Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at www.linuxforu.com. The sender of each published tip will get an LFY T-shirt.

Fine Tuning the Environment

...and Making Passes at the Command Line

If you know a few 'idioms', you can control and fine-tune a process' environment.

Since we breathe in the air in our environment, it makes sense for us to improve the quality of that air by using less fossil fuels and planting more trees. Similarly, when Linux processes are born, they start interacting with an inherited environment. There are some pretty cool tools that Linux provides to inspect this environment. By learning a few of these, we can control and fine-tune the process' environment.

Many important variables reside in the environment. `PATH` is one and `LD_LIBRARY_PATH` is another. At times, environment variables affect the execution of a process. They are like control knobs, using which you can change settings. For example, the manual (`man`) for `malloc` mentions the environment variable `MALLOC_CHECK_` which, when set, substitutes another implementation of the allocator, and catches memory corruptions in C code.

Okay, so that's enough of talking—let us get the ball rolling! We know a Linux process is a running program. Creating a process is simple:

```
~/atul> cat > ./tst.sh
sleep 7200
~/atul> chmod +x !$
chmod +x tst.sh
~/atul> ./tst.sh & # start it in background
[1] 32069
~/atul>
```

Let us inspect the environment of this running process we created. Here, the PID of the process created from the shell script is 32069.

```
~/atul> ps eeewww -p 32069
PID TTY STAT TIME COMMAND
```

```
30988 pts/1 S+ 0:00 sh ./tst.sh HOSTNAME=hydqa25.in.ibm.com
TERM=xterm SHELL=/bin/bash...
~/atul>
```

Oops, this is quite a listing... But hey, there's no need to rub our eyes—a little scripting can make things pretty palatable.

```
~/atul> ps eeewww -p 32069 | tr '[:space:]' '\n' | grep '='
HOSTNAME=hydqa25.in.ibm.com
TERM=xterm
SHELL=/bin/bash
...
~/atul>
```

Every assignment you see is an environment variable.

If the environment listing is too long, we can always pass it along to `less` and look at it in a *more* leisurely way.

So there's no more guesswork when you suspect some environment variable is not set correctly—just ask `ps` to probe the suspicious process and check out the environment.

Important aside (a.k.a Idiom 1)

Shells also have the following rule: any variable assignment *preceding the command name* turns the variable into an environment variable. So, for example,

```
~/atul> a=1 b=2 c=3 command arg1 arg2 ...
```

...creates three environment variables `a`, `b` and `c` and puts these into the environment of the process created as a result of running a command.

We use the idiom given below to fine tune the environment.

Wiping the slate clean (a.k.a Idiom 2)

At times, it is a good idea to control the environment variables—while isolating problems, you want to be absolutely sure the environment is pristine.

Here is how to do it:

```
~/atul> env -i ./tst.sh &
[1] 24302
~/atul> ps eeewww -p 24302
PID TTY STAT TIME COMMAND
24302 pts/1 SN 0:00 /bin/sh ./tst.sh
~/atul>
```

Now, I am sure I am not inheriting any environment variables.

Fine tuning

Create a file with your chosen environment variables, one assignment, per line:

```
~/atul> cat tst.env
x=25
y=26
~/atul> env -i $(cat tst.env) ./tst.sh &
[2] 24618
~/atul> ps eeewww -p 24618
PID TTY STAT TIME COMMAND
24618 pts/1 SN 0:00 /bin/sh ./tst.sh x=25 y=26
~/atul>
```

It is interesting and enlightening how the above works. All shells parse the command line. During this parsing process, they also perform various operations—one of which is substitution.

Here we use process substitution to run the *cat* command. *cat* displays the contents of the *tst.env* file to its standard output. This output is captured by the shell and split into words—the *word splitting* operation.

So, in this case, every assignment is a word, and the shell processing tacks these before the command name, effectively creating those environment variables of your choice.

Oh okay, but what about any newlines in the *tst.env* file? For the word splitting operation, both newline and spaces are delimiters (as are tabs). To see for yourself:

```
~/atul> echo $(cat tst.env)
```

IFS has the default value here—space, tab, and newline—versus:

```
~/atul> OLDIFS=$IFS
~/atul> IFS=' '
~/atul> echo $(cat tst.env)
~/atul> IFS=$OLDIFS
```

The first version gives everything on the same line,

while the second preserves the newlines, as we removed the newline char from IFS. The newlines no longer act as the delimiter.

Aha! So we have some fine tuning to be done.

However, what about embedded spaces in the value of our environment variables?

```
~/atul> cat >> tst.env
z="this is a variable with a space and tab"
^D
~/atul> env -i $(cat tst.env) ./tst.sh
env: is: No such file or directory
~/atul>
```

The problem is, that by the time we finished *word splitting*, the " (double quotation mark) is just another character—it has lost its meaning as a protector of spaces. To understand this better, I would advise going through the shell's man pages to understand the flow. A flow chart showing how a shell processes its command line would also help.

Making passes at the command line (a.k.a the eval magic; or Idiom 3)

Don's *Art Of Computer Programming* has the following joke:

Old lady on the bus: "Young boy, can you tell me how to get off at Pasadena Street?"

Young boy: "Just watch me, and get off two stops before I do."

The joke is that the boy gives a *two pass* algorithm.

To solve the above problem we need to do one more pass before we hand over the command to *env*. Here is how to do it—just stick an *eval* before *env*.

```
~/atul> eval env -i $(cat tst.env) ./tst.sh
```

The first pass pulls the assignment strings out of the file and lays these before the command. The second pass then correctly parses these assignments and executes the command—this preserves the significance of the quotes.

Open another terminal and check out the environment. You will see the correct value of *z*. Another way to check is to stick a *echo "\$z"* in *tst.sh* before the *sleep* command and see for yourself.

Linux is so environment friendly! 

By: Atul Shrinivas Khot

The author works at IBM Software Labs in Hyderabad as a senior software engineer, and has been dabbling with UNIX/Linux for the last 14 years. He is into Java/J2EE, Groovy and Ruby these days, but at times he hankers after dear-old C++ and Perl. He loves design patterns, algorithms, multi-threading and refactoring code to make it stylish. And of course, he loves vim (never miss a chance to use it) and zsh. He collects classic British mysteries (especially those green and white Penguin crime classics—penguins make life so delectable;-))



Learn how to tune your hard disk I/O using `read_ahead`, elevators, queuing theory and the utilisation law.

Performance Tuning and Monitoring

Last month we learned how to check the capabilities of our existing hardware, and how to ascertain the current load on the system using tools like *iostat*, *dmidecode*, *sar*, *mrtg*, *gnuplot*, etc. Well, it's the right time to tune your system to your needs, after monitoring the data now available to you.

First, tune your HDD I/O.

Tuning disk I/O using elevators

Hard disks are electro-mechanical devices with two kinds of components—internal and external. While the internal component is essentially the sealed chamber called the hard drive assembly (HDA), the external components are located on a printed circuit board (PCB), often referred as the logic board. Refer to Figures 1 and 2.

Today's hard disk drives are made of platters coated with magnetic material. An electromechanically actuated armature contains the 'read/write' heads that move between these platters. The armature is normally composed of multiple arms. All the arms are bound together to work as a single unit. So when the armature moves to a specific track on the platter, all the read/write heads are placed on the same track on each platter. This is known as a 'cylinder'.

Remember that whenever there is a moving part involved in I/O, there will be latency. As our

HDD also has moving parts—the platters and the armature—we cannot expect a 'wire speed' I/O. So whenever the disk controller requests a read/write from a particular sector, the armature has to move from its current position to the new requested position. This delay is known as the seek time. Add the time the motor takes to make that particular sector appear under the read/write head to the seek time—this time is called the rotational delay, or 'rotational latency'.

The good news is that all modern HDDs (more than 90 per cent of all known brands) use a concept called ZCAV (Zonal Constant Angular Velocity). This takes advantage of the fact that more linear space is available on the outer tracks of the disk platter rather than on the inside tracks. Now since the disk spins at a constant speed, which is also known as CAV (Constant Angular Velocity), the read/write I/O speed will be greater at the outer tracks as compared to the inner tracks. You can use a program called Bonnie++ [www.coker.com.au/bonnie++] to check your hard disk.

So the general rule is that you should always create the locations that need frequent I/O—for example, `/home` and `swap`—on the outer tracks. Since generally all HDDs allot partitions from the outside, the easiest way to achieve this is to create these partitions first when partitioning your hard disk.

BUS considerations

Most PC-based systems use the PCI (Peripheral Component Interconnect) bus to connect devices to the processor. The performance of the PCI is determined by the size of data path (bus width) and the speed at which it transfers the data (clock rate). Currently, the PCI supports the 32/64-bit bus width at a speed of 66 MHz.

There is also a parallel interface standard called SCSI (Small Computer System Interface). A typical problem that arises with the SCSI occurs when you place devices of different speeds on the same bus. So a device with a faster throughput may get 'stepped down' to the speed of the slower device on the bus. You can use the *sginfo* tool, part of the *sg3_utils* package, to query or change the parameters of SCSI devices on your machine.

Tuning the sequential read access

The kernel, when reading a file, always tries to take advantage of sequential disk access. 'Read-ahead' is based on the same assumption that if an application is accessing data from Block A, then the chances are more likely that the next Blocks—B, C and D will also need to be read. Therefore, by actually reading ahead and then caching the pages in memory, the kernel improves the I/O and reduces the response time. However, the read-ahead algorithm is designed to get turned off automatically whenever a random read request is detected.

The read-ahead function is based on two values:

1. The current window—the application reads pages from the buffer cache that is its current window.
2. The 'ahead' window—while the application reads pages from the current window, IO happens on the 'ahead' window.

When the application has finished reading from the current window, the 'ahead' window becomes the current window and a new 'ahead' window is started.

To view the current window size, issue `cat > /sys/block/sda/queue/read_ahead_kb`. The following is the command output for my machine:

```
# cat /sys/block/sda/queue/read_ahead_kb
128
```

As you can see, the current and default read-ahead is 128 KB. I can tune it to 256 KB, for example, as follows:

```
# echo 256 > /sys/block/sda/queue/read_ahead_kb
# cat /sys/block/sda/queue/read_ahead_kb
256
```

As is evident, the current window has now become 256 KB. However, note that rebooting the machine will change the default value to 128 KB. Please feel free to use the */etc/rc.local* file to make your changes permanent.

We can also use the *blockdev* command to report read-ahead in sectors. A sector is always 512 bytes in kernel 2.6.x.

```
# blockdev --report /dev/sda
```




Figure 1: A typical hard disk drive




Figure 2: Components of the hard disk

```
RO RA SSZ BSZ StartSec Size Device
rw 256 512 4096 0 320072933376 /dev/sda
```

Remember that if page access occurs in the current window then the size of the new read-ahead windows increases by two pages. So because the default current window's size is 128 KB, the current read-ahead appears as 256. We can use the *blockdev* command with the following options to view and then set our read-ahead.

```
# blockdev --getra /dev/sda
256
# blockdev --setra 512 /dev/sda
# blockdev --getra /dev/sda
512
```

Tuning the disk queue

The I/O subsystem is a series of processes with the responsibility to move blocks of data between the hard disk and the RAM. Generally, every computer task consists of a utility performing either one or both of the following:

- Read a block of data off the disk and move it to RAM
- Write a new block of data from the RAM to disk

These read/write requests are transformed into 'block device requests' that go into a queue. When adding an entry to

the queue, the kernel first tries to enlarge an existing request by merging the new request with one that is already in the queue. If the request cannot be merged with an existing one, then the new request will be assigned a position in the queue based on several factors including the elevator algorithm (which we will configure later).

First, let's check what the current queue length is and which current IO scheduler the OS is using. To view the current queue length, issue the following command:

```
# cat /sys/block/sda/queue/nr_requests
128
```

You can see that my current queue length is 128 requests. Now I want to increase it to 170—let's do it using the *echo*:

```
# echo 170 > /sys/block/sda/queue/nr_requests
# cat /sys/block/sda/queue/nr_requests
170
```

To make it permanent, add the entry to the */etc/rc.local* file. [As a rule, keep in mind that anything under */sys* can be made permanent by using */etc/rc.local*.]

Now, let's check which elevator (or I/O scheduler) algorithm the OS is using by default:

```
# cat /sys/block/sda/queue/scheduler
noop anticipatory deadline [cfq]
```

On my system, it's cfq. Now the question is...

What is an I/O scheduler?

I/O schedulers control the way Linux handles reads and writes to disks. Linux 2.6 includes a number of I/O schedulers—the intention of providing these choices is to allow better optimisation for different classes of workload.

Without an I/O scheduler, Linux follows FIFO (First In, First Out). This could result in massive HDD thrashing—for example, if one process is reading from one part of the disk and one writing to another, the heads would have to seek back and forth across the disk for every operation. The scheduler's main goal is to optimise disk access times.

This is where I/O schedulers come into the picture. Quoting the WlugWiki [www.wlug.org.nz/LinuxIoScheduler] an I/O scheduler can use the following techniques to improve performance:

- Request merging—The scheduler merges adjacent requests together to reduce disk seeking.
- Elevator—The scheduler orders requests based on their physical location on the block device, and it basically tries to seek in one direction as much as possible.
- Prioritisation—The scheduler has complete control over how it prioritises requests, and can do so in a number of ways

In addition to these, all I/O schedulers also take into account resource starvation, which ensures requests to

be serviced.

Scheduler algorithms operate like elevators in buildings, and are therefore also called elevators. As per *The Red Hat Enterprise Linux 5 IO Tuning Guide*: “The algorithms used to operate real-life building elevators make sure that it services requests per floor, efficiently. To be efficient, the elevator does not travel to each floor depending on which one issued a request to go up or down, first. Instead, it moves in one direction at a time, taking as many requests as it can until it reaches the highest or lowest floor, then does the same in the opposite direction. Simply put, these algorithms schedule disk I/O requests according to which logical block address on the disk they are targeted to. This is because the most efficient way to access the disk is to keep the access pattern as sequential (i.e., moving in one direction) as possible. Sequential, in this case, means 'by increasing the logical block address number'. As such, a disk I/O request targeted for disk block 100 will normally be scheduled before a disk I/O request targeted for disk block 200. This is typically the case, even if the disk I/O request for disk block 200 was issued first.”

There are currently four types of schedulers available:

1. The no-op scheduler (noop)
2. The anticipatory IO scheduler (anticipatory)
3. The deadline scheduler (deadline)
4. The Complete Fair Queueing scheduler (CFQ)

The no-op scheduler

As the name suggests, ‘no-op’ actually does nothing—it simply works in the FIFO manner. No configurable options are available here. The I/O requests are sent in the same manner as they are received, and it is left to the hardware to work on it. The main goal is to conserve CPU cycles.

You can make no-op your IO scheduler as follows:

```
echo noop > /sys/block/sda/queue/scheduler
```

To make it your choice of I/O scheduler across reboots, edit */boot/grub/grub.conf* and append *elevator=noop* at the end of the kernel line:

```
root (hd0,0)
kernel /vmlinuz-2.6.29.6-217.2.16.fc11.i686.PAE ro root=LABEL=Fedora rhgb quiet
elevator=noop
```

The anticipatory scheduler

I personally think this is the most optimistic of the schedulers. It waits in anticipation that the next I/O request might be for the next disk block, before moving to another location. The basic idea is that the IO scheduler will wait for a certain short period of time after catering to one I/O request and before going on to the next I/O request. So if the next I/O request is for the next disk block, it saves time in moving the head back to the same location again. But this may also result in additional latency, if the next I/O is not for the next disk block.

Switch to this scheduler as follows:


```
# echo anticipatory > /sys/block/sda/queue/scheduler
```

Add *elevator=anticipatory* in the kernel line of */boot/grub/grub.conf*, if you want to make it your default.

But how can you tune it as per your needs? Well, the primary tunables are under the */sys/block/sda/queue/iosched/* directory. Some common values that can be tuned are (the values appear automatically under */sys/block/sda/queue/iosched/* as soon as you change your elevator):

- **antic_expire** – Here you can set the maximum amount of time (in milliseconds) you anticipate a good read (one with a short seek distance from the most recently completed request) before giving up. For example, I can change the *antic_expire* from the default 6 ms to 60 ms, as follows:

```
# cat /sys/block/sda/queue/iosched/antic_expire
6
# echo 60 > /sys/block/sda/queue/iosched/antic_expire
# cat /sys/block/sda/queue/iosched/antic_expire
60
```

You can make it permanent by adding the echo line in your */etc/rc.local* file.

- **read_expire** – All the read and write IO requests are processed in batches. So this parameter specifies the time within which the read request must be fulfilled or completed before it expires. The time is specified in milliseconds.

```
# cat /sys/block/sda/queue/iosched/read_expire
125
# echo 250 > /sys/block/sda/queue/iosched/read_expire
# cat /sys/block/sda/queue/iosched/antic_expire
250
```

- **write_expire** – they are equivalent to the above, for writes.

```
# cat /sys/block/sda/queue/iosched/write_expire
250
# echo 300 > /sys/block/sda/queue/iosched/read_expire
# cat /sys/block/sda/queue/iosched/antic_expire
300
```

The deadline scheduler

As the name suggests, here all the I/O requests have to be fulfilled in a specified amount of time—before they expire. So when an I/O request enters the queue, it is assigned some time (in milliseconds), within which it has to be fulfilled before the time expires, regardless of its targeted block device. The same applies for the write requests. It will also specify how the maximum number of read requests will be fulfilled before switching to a write request.

In order to switch to this scheduler, use the following command:

```
# echo deadline > /sys/block/sda/queue/scheduler
```

To make it permanent, add *elevator=deadline* in the kernel line of */boot/grub/grub.conf*.

And here are its tunables (located under the same */sys/block/sda/queue/iosched/* directory):

- **read_expire** – Here, when a read request first enters the I/O scheduler, it is assigned a deadline that is the current time plus the *read_expire* value, in milliseconds. And, like before, I can change the default value for *read_expire* from 500 ms to 444 ms:

```
# cat /sys/block/sda/queue/iosched/read_expire
500
# echo 444 > /sys/block/sda/queue/iosched/read_expire
# cat /sys/block/sda/queue/iosched/antic_expire
444
```

- **write_expire** – As you might have already guessed, its function is the same as *read_expire*, but it is for writes. The time is again in milliseconds.

```
# cat /sys/block/sda/queue/iosched/write_expire
5000
# echo 4000 > /sys/block/sda/queue/iosched/read_expire
# cat /sys/block/sda/queue/iosched/antic_expire
4000
```

- **front_merges** – Normally, I/O requests are merged at the bottom of the request queue. This Boolean value controls whether an attempt should be made to merge the request to the top of the request queue. A value of 0 indicates that the front merges are disabled.

```
# cat /sys/block/sda/queue/iosched/front_merges
1
```

The Completely Fair Queue scheduler (CFQ)

The CFQ scheduler works in a pure, democratic way. It equally divides all the available bandwidth between all the I/O requests. CFQ uses 64 internal queues to maintain and keep I/O requests. It fills internal queues in a round-robin manner and pulls the I/O requests from these 64 internal queues and places them into a dispatch queue, where they are catered to. Internally, CFQ also changes the order of I/O requests to minimise the head movement.

Most Linux distros come with CFQ as the default I/O scheduler. If not, make it the default as follows:

```
# echo cfq > /sys/block/sda/queue/scheduler
```

Append *elevator=cfq* to the kernel line of */boot/grub/grub.conf* to make the change permanent.

As for the tunables, they are:

- **quantum** – This is the total number of requests to be moved from internal queues to the dispatch queue in each cycle.

```
# cat /sys/block/sda/queue/iosched/quantum
4
# echo 8 > /sys/block/sda/queue/iosched/read_expire
# cat /sys/block/sda/queue/iosched/antic_expire
8
```

Here I just changed the number of requests to 8 from the default value of 4. Feel free to use */etc/rc.local* to make the changes permanent.

- **queued** – This is the maximum number of requests allowed per internal queue. You can view and change the current value using the *cat* and *echo* commands, respectively, as above.

An introduction to queuing theory

Quoting Wikipedia on the queueing theory, Little's Law states: "The long-term average number of customers in a stable system L is equal to the long-term average arrival rate, λ , multiplied by the long-term average time a customer spends in the system, W , or: $L = \lambda W$. Here:

- L = queue length, i.e., the average number of requests waiting in the system
- A = arrival rate, which is the rate at which requests enter a system
- W = wait time, which is the average time to satisfy a request

At a constant arrival rate (λ), if the queue length (L) increases, so will the wait time (W), and if the queue length (L) decreases, so will the wait time (W). So at a constant (λ), L is directly proportional to W .

Wait time (W) consists of queue time (time waiting for a resource to become available, Q) and service time (time for a resource to process a request, S). So we can say, $W = Q + S$. So Little's Law can be stated as: $L = \lambda (Q + S)$

The Utilisation Law is a derivative of Little's Law. It simply states that: $U = (\text{service time}) \times (\text{arrival rate})$. For a saturated resource U is 1. This means, if the value of U reaches near 1 (or becomes 1), the resource cannot take any more requests.

Now let's check the utilisation of my server hard disk. Run `iostat -x -d /dev/sda 1` to get the values of the average reads per second, the average writes per second and the average service time. I get the following output:

```
Device: rrqm/s wrqm/s r/s w/s rsec/s wsec/s avgrq-sz avgqu-
sz await svctm %util
sda      48.00 0.00 86.00 0.00 17760.00 0.00 206.51 0.27
3.12 2.84 24.40
```

From this, I get:

- Average read request (`avgrq-sz`) = 206.51
- Average write request (`wsec/s`) = 0
- Average service time (`svctm`) = 2.84 ms

Now we know that $U = S \times A$ (can be remembered as USA, :-)). Hence, $U = [(206.51 + 0) \times (2.84)] / 1000 = 0.586$. So my utilisation is just half way to the saturation level of 1. If this number gets closer to 1 then it's time to change the disk.

We can also find out how many requests per second this hard disk can take. Just make $U=1$ and find out the A (arrival rate). That is, if $1 = 2.84\text{ms} \times A$, then $A = 1000/2.84 = 352.112$ requests per second. So this hard disk can take a maximum of 353 requests per second. Anything over this number can cause the machine to fall over.


Finding hot spots in the code

We can use the `strace` tool to trace system calls and signals. `strace` runs a specified command until it exits or ends. It intercepts and records the system calls made by the process and the signals that are received by the process.

For example, let's find out the system calls made by the `who` command:

```
# strace -c who
root pts/1 2009-09-13 10:37 (172.24.0.123)
alok :0 2009-09-13 10:38
alok pts/0 2009-09-13 10:38
alok pts/1 2009-09-14 00:16
% time seconds usecs/call calls errors syscall
-----
nan 0.000000 0 20 read
nan 0.000000 0 4 write
nan 0.000000 0 20 open
nan 0.000000 0 21 close
nan 0.000000 0 1 execve
nan 0.000000 0 1 time
nan 0.000000 0 45 alarm
nan 0.000000 0 2 2 access
nan 0.000000 0 4 kill
nan 0.000000 0 3 brk
nan 0.000000 0 4 munmap
nan 0.000000 0 3 mprotect
nan 0.000000 0 2 _llseek
nan 0.000000 0 30 rt_sigaction
nan 0.000000 0 22 mmap2
nan 0.000000 0 7 1 stat64
nan 0.000000 0 20 fstat64
nan 0.000000 0 31 fcntl64
nan 0.000000 0 1 set_thread_area
-----
100.00 0.000000 241 4 total
```

As you can see, it clearly tells you the name of the system call and how many times it had been called.

Here I conclude for this month. Next time we will learn how to configure a tuneable RAID, using the external journal, and learn about processor scheduling using a locality of reference. **END** 

References and attribution

- RHEL 5 IO Tuning Guide: <http://www.redhat.com/docs/wp/performancetuning/iotuning/index.html>
- Linux I/O Scheduler [WlugWiki]: www.wlug.org.nz/LinuxIOScheduler

By: Alok Srivastava

The author is the founder of Network NUTS and holds MCP, MCSE, MCDBA, MCT, CCNA, CCNP, RHCE and RHCSS certifications. Under his leadership, Network NUTS has been a winner of the "Best Red Hat Training Partner in North India" for the last three years in a row. He has also been a trainer for the Indian Air Force, NIC, LIC, IFFCO, Wartsila India Ltd, the government of Rajasthan, Bajaj Allianz, etc. You can reach him at alok@networknuts.net.

South Delhi centre is the India's only Dedicated Learning Centre for Red Hat Linux.



for networking, for life

RHCA*

RHCSS

RHCE

Open LDAP

Virtualization

Shell Scripting

"RH436
(Red Hat Enterprise
Clustering and Storage
Management) batch
starting from
12th-15th Nov.09"

Congratulations



Ms. Mandeep Kaur

on achieving highest ever aggregate
of 99.5% in RHCSS (Ex333)*



Boot Camps - RHS423 : 20th Oct. - 23rd Oct.09

RHS333 : 24th Oct. - 27th Oct.09

RHS429 : 28th Oct. - 31st Oct.09

EX333 14.10.09 EX423 13.10.09 EX429 15.10.09

RHCE 12th & 30th Oct, 09

Fact #1

Given 1783+ Red Hat Certified professional
(RHCE + RHCSS), with BEST result in
Asia-Pacific region.

redhat
PREMIER
PARTNER



South Delhi : A-184, Bhisham Pitamah Marg, Opp. Defence Colony, Delhi-2 Tel. : 46526981-2, 9310024503

East Delhi : 508, DDA Complex, Laxmi Nagar Distt. Centre, Delhi-92 Tel. : 42448951-2, 9310024502

North Delhi : 2513, Hudson Lane, Kingsway, Camp, North Campus, Delhi-9 Tel. : 65156734, 9310024501

www.networknuts.net

* among girls.



Sandya Mannarswamy

Welcome to this month's CodeSport, in which we continue our discussion on multi-threaded code and discuss lock granularity.

Last month's takeaway problem came from one of our readers, Nilesch Govande. Given a two-dimensional $n \times n$ matrix M of positive and negative integers, readers were to find the sub-rectangle with the largest sum. The sum of a rectangle is the sum of all the elements in that rectangle.

In this problem, the sub-rectangle with the largest sum is referred to as the *maximal sub-rectangle*. A sub-rectangle is any contiguous sub-array of size 1×1 or greater, located within the whole $n \times n$ array.

Three of our readers had sent their answers to this problem. However, none of them were correct. Hence I am going to keep this problem open for this month also. However, let me provide some clues.

The problem of a maximal sub-rectangle is based on a similar variant of the problem, namely that of calculating the maximum contiguous sub-array in one dimension.

That problem can be solved in $O(n)$ time (time linear to the size of the input array). We had discussed this problem in an earlier column. Therefore, extending that idea, the first approach would be to compute the sums of all possible rectangular sub-arrays $A[i1-i2][j1-j2]$

where $0 < i1, i2 < n$ and $0 < j1, j2 < n$.

Let $S[r,c]$ denote the sum of the rectangular sub-array of M with one corner at entry $[0,0]$ and the other at $[r, c]$. This can be computed in $O(n^2)$ time.

Observe that the sum of any rectangular sub-array of M can be computed in constant time given the table S . Once we have the table S , we can just guess different values for the lower-left and the upper-right corner of the rectangular sub-array and use the S table to compute its sum. The sub-array that yields the maximum sum is our result. Such an algorithm will be of the order of $O(n^4)$.

I leave it to the readers to write down the actual code to compute this. Can you come up with a $O(n^3)$ algorithm? Please do send in your solutions.

Now we will continue our discussion on multi-threaded code, while looking at locks and the issue of lock contention.

Mutual exclusion locks

As we all know, a binary semaphore or a mutex, as it is more popularly known, is the simplest type of lock. Though there are other types of locks such as adaptive locks, read-write locks, queue-based locks, etc, for our discussion on lock contention, we will focus on the mutex.

If there is a critical section of code that can be executed by only one thread at a time, a mutex is associated with the critical section. Before a thread enters the critical section, it needs to acquire the lock associated with that critical section. If the lock is already held by another thread, then the current thread needs to wait until the currently executing thread releases the lock.

Let T1 be the thread requesting the lock and let T2 be the thread currently holding the lock. What should T1 do, when it finds that the lock is not available? Should it keep checking the lock variable until it becomes free, or should it just put itself to sleep for a short duration and then come back and check whether the lock has been released? Depending on the nature of the application, either option can be beneficial.

A spinlock is a lock in which the thread simply waits ('spins') until the lock becomes available. It is very efficient if the threads are only likely to be blocked for a short period of time, as it avoids the overhead of operating system process re-scheduling.

It is wasteful if the lock is held for a long period of time. It is also possible for the thread to spin for a short duration on the lock and if the lock is not released during that duration, it will go to sleep and get rescheduled by the OS. Consider the following code snippet where the thread continues to spin on the lock until it acquires the lock:

```
pthread_mutex_t mutex1;
int all_things_done;

void my_thread_work(void)
{
    RETRY:  if (pthread_mutex_lock(&mutex1))
    {
        goto RETRY;;
    }
    Things_done++;
    Pthread_mutex_unlock(&mutex1);
}
```

This is an example for a pure spinlocking code, where the thread requesting the lock continues to spin on the CPU until it gets the lock or until its CPU quantum is complete. If the critical section protected by the spinlock is long, then the requesting thread unnecessarily ends up wasting CPU cycles by spinning.

On the other hand, if the critical section protected by the spinlock is quite short, then the above code snippet is preferable, since the thread can acquire the lock instead of having to be rescheduled by the OS.

Depending on the nature of the critical section, a programmer can decide to use a pure spinlock or a lock, which spins for a short duration and then goes off to sleep if the lock is not obtained during the spin duration.

```
pthread_mutex_t mutex1;
int all_things_done;
void my_thread_work(void)
{
    RETRY:  if (pthread_mutex_lock(&mutex1))
    {
        retry_Count++;
        If (retry_count < MAX_RETRY_COUNT)
            goto RETRY;;
        sleep(SLEEP_DUR);
        goto RETRY;
    }
    Things_done++;
    Pthread_mutex_unlock(&mutex1);
}
```

In the above code snippet, the requesting thread spins until the number of tries it makes to acquire the lock reaches the MAX_RETRY_COUNT. After that, it goes to sleep and can be rescheduled by the OS. The values of MAX_RETRY_COUNT and SLEEP_DUR are, typically, application and critical section dependent.

Lock granularity

An important property of a lock is its granularity. The granularity is a measure of the amount of data the lock is protecting. In general, choosing a coarse granularity (a lock protecting a large segment of data) results in less lock overhead when few threads are accessing the protected data, but worse performance when multiple threads are running concurrently and accessing different parts of the protected data. This is because of increased lock contention.

Lock contention occurs whenever one thread attempts to acquire a lock held by another thread. The more granular the available locks, the less likely it is that one thread will request a lock held by the other, since the data areas they are accessing are different.

The more coarse the lock, the higher the likelihood that the lock will stop an unrelated thread from proceeding. Conversely, using a fine granularity (a larger number of locks, each protecting a fairly small amount of data) increases the overhead of the locks themselves but reduces lock contention.

More locks also increase the risk of deadlock. Consider the following code snippet, which shows the skeleton code for ticket booking in a cinema theatre. The two-dimensional matrix 'tickets' represents the tickets available in a cinema hall. The seating in the cinema hall consists of rows and columns. A ticket is represented by a unique tuple of (row, column).

Ticket booking is done by a multi-threaded application, with each thread representing a booking window counter. The cinema hall has multiple booking windows. Hence, access to the matrix 'tickets' needs to be protected by a lock and performed inside

a critical section. Consider the code snippet below:

```
int tickets[MAX_ROWS][MAX_COLS];
pthread_mutex_t tickets_lock;

void BookTicket(int row, int column)
{
    pthread_mutex_lock(&tickets_lock);
    tickets[row][column].status = 'booked';
    pthread_mutex_unlock(&tickets_lock);
}
```

Here, we have a single mutex lock *tickets_lock* that protects the entire tickets table. At a time, only one thread can hold the lock and hence only one ticket can be booked, even though there can be two people requesting tickets such as (1,1) and (10,10) which are totally unrelated tickets. Since we use a single lock to protect the entire tickets table, we prevent concurrency among independent accesses to unrelated parts of the tickets table.

The above code snippet is an example of coarse-grained locking. Now we can reduce the granularity of the lock, by associating a unique lock with each row and each column in the cinema hall. Consider the code snippet given below:

```
int tickets[MAX_ROWS][MAX_COLS];
pthread_mutex_t row_lock;
pthread_mutex_t column_lock;

void BookTicket(int row, int column)
{
    pthread_mutex_lock(&row_lock);
    pthread_mutex_lock(&column_lock);
    tickets[row][column].status = 'booked';
    pthread_mutex_unlock(&column_lock);
    pthread_mutex_unlock(&row_lock);
}
```

In this snippet, we can allow two threads that are trying to book tickets (i1,j1) and (i2,j2) such that i1 is not equal to i2 and j1 is not equal to j2, to book tickets concurrently. This increases the concurrency compared to the coarse grained locking case when we used a single lock for the whole tickets table.

Still, there is concurrency to be gained by using finer grained locks, since a ticket is uniquely specified by a row and column, and as long as the two threads access two tickets in separate locations [row1][column1] and [row2][column2], where either 'row1' is not equal to 'row2' or 'column1' is not equal to 'column2', the threads can book the tickets concurrently.

In order to allow this, we need to associate a lock with each unique seat location, namely, each element of the tickets matrix. This is demonstrated in the code snippet below:

```
int tickets[MAX_ROWS][MAX_COLS];
pthread_mutex_t ticket_locks[MAX_ROWS][MAX_COLS];

void BookTicket(int row, int column)
{
    pthread_mutex_lock(&ticket_locks[row][column]);
    tickets[row][column].status = 'booked';
    pthread_mutex_unlock(&ticket_locks[row][column]);
}
```

However, note that we end up using as many mutex locks as there are elements in the tickets matrix. This can lead to excessive lock overhead, since for each 'tickets' element access, we have to access one mutex lock.


Earlier, when we used coarse grained locks, a single mutex was used to protect multiple 'tickets' elements. This resulted in lowered concurrency, but also reduced the locking overhead.

Hence, the lock granularity one needs to choose for an application is often dependent on the access patterns of the data shared by the threads. The locking overhead associated with fine grained locks needs to be traded off against the increased concurrency achieved by fine grained locking.

Note that lock granularity has a major impact on lock contention, which we shall discuss in our next column.

Takeaway problem for this month

As I mentioned before, we keep open the takeaway problem from last month since there were no correct solutions. Based on the hints given in this column, please do send me your solutions to the maximal sub-rectangle problem.

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me, along with solutions and feedback to sandyasm_AT_yahoo_DOT_com. Till we meet again next month, happy programming!  **END**

References

- [en.wikipedia.org/wiki/Lock_\(computer_science\)](https://en.wikipedia.org/wiki/Lock_(computer_science))
- en.wikipedia.org/wiki/Spinlock

About the author:

Sandya Mannarswamy. The author is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.

Your Daily Dose of **TECHNOLOGY NEWS**



IF IT'S TECHNOLOGY IT'S HERE

**Free
Subscription**

Simply visit www.efytimes.com
and **subscribe** to the newsletter

DAILY TECH UPDATES >>>

From the EFY Group on subjects ranging from nanotechnology to open source software can land in your mailbox for **FREE!**

Hardware, Software, Consumer Electronics, Robotics, Automotives, Telecom, Space—any news that's related to technology will be featured here.

www.efytimes.com



EFY Enterprises Pvt Ltd

D-87/1, Okhla Industrial Area, Phase I, New Delhi 110 020; Ph: +91-11-26810601-03
Fax: +91-11-26817563, 26812312; E-mail: info@efyindia.com; Website: www.efyindia.com

OpenSolaris on Xen **Part 1**

The Ecosystem of the Xen Platform

This article deals with the concepts associated with Xen virtualisation technology—something we should be aware of before we get started with Xen deployments.

Being able to run multiple operating systems concurrently on the same computer is gaining currency, of late, since this reduces costs and increases utilisation of systems.

Being able to carve multiple virtual machines (VM) out of a single physical machine is called virtualisation.

Virtualisation of the underlying platform can be achieved in multiple ways. For example, VirtualBox takes an unmodified OS and runs it as a process inside another OS. Another approach is to run a thin layer of software called a hypervisor on the bare metal and make all the guests run on top of the hypervisor. Xen adopts this model.

It uses para-virtualisation (PV) to boost the performance of the VMs. To run as a PV guest, an OS needs to be ported to the Xen architecture—which is very similar to the actual underlying platform (i.e., the x86/x64 architecture). Thus, a PV guest is aware that it is running on top of a hypervisor and hence not in control of the underlying physical machine.

In this four-article series we shall explore: 1) The ecosystem of an OpenSolaris Xen platform and other under-the-hood machinations of the hypervisor; 2) OpenSolaris as dom0; 3) The changing architecture of Xen in the face of new hardware advances; and 4) A high level code walk-through of the dom0 kernel.

The ecosystem of a Xen platform

dom0: The first VM that comes up on a Xen system is referred to as dom0, and it has extra privileges compared to other VMs that come up later. It provides all the services mentioned below.

XenStore: Configuration information, such as which block/network devices are associated with a domU, is written to XenStore by dom0, which is, in turn, used by domU to boot up [similar to what the BIOS provides for an OS that comes up on metal]. It also provides information about currently active domains. Dynamic reconfiguration of guest domains is made possible by XenStore. For example, you can pass additional disk devices from dom0 to domU, while domU is running.

xend: This provides the interface between the hypervisor and the rest of the user-space tools. It also provides a stable API for the higher-level management tools to be built without worrying about hypervisor changes. xend performs most of the work involved in a guest domain's lifecycle—creation, shutdown, suspension and migration, in response to admin commands from dom0. It writes the resource information associated with a guest to XenStore, besides listening to a Unix socket as well as httpd—hence making local as well as remote management of a Xen system, possible.

XenBus: A guest kernel listens on a XenBus, waiting for any events that will result in dynamic reconfiguration of the guest. XenStore delivers those events from dom0.

Pygrub: It understands the guest filesystems, and reads the kernel binary, boot archive binary and any other module required to boot up a guest system.

libvirt and tools: As more virtualisation technologies are hitting the market, high-level management tools should be immune to the underlying hypervisor. The libvirt library provides such hypervisor agnostic APIs to build higher-level management tools. virt-install (1) is one such tool, built on top of libvirt to install a guest. virsh (1), the basic administrative tool to interact with the Xen system, is also built on top of libvirt.

vdiskadm: It allows us to create virtual disk images of various formats—vmdk (VMware format), vdi (VirtualBox format), vhd (Hyper-V) and raw. You can take snapshots and make clones of vdisk images for ease of management. virt-convert (1) allows conversion of one disk format to another disk format. By default, OpenSolaris uses the vmdk format.

vdisk: For each virtual disk passed to a guest domain, there is one vdisk process running in the user land of the dom0, to coordinate the IO. This prevents proliferation of 'disk format' knowledge in the hypervisor or dom0 kernel.

libvirtd: On an OpenSolaris-based Xen system,

management tools communicate with xend via libvirtd to improve security.

Log files: Detailed logs of hypervisor activities and domain life cycle events are available in the `/var/log/xen/` directory.

Machinations of the Xen hypervisor

Hypercall interface: In the same way applications make use of system calls to request the services of the kernel, a guest kernel makes hypercalls to request the services of the hypervisor. Applications running on top of guest kernels continue to run unmodified.

Virtualising the CPU: x86 CPUs provide four privilege levels—0 to 3, with 0 being the highest privilege level and 3 being the lowest. Because of the multiple rings of protection, it is easy to virtualise x86-based CPUs. In 32-bit mode, Xen runs in Ring 0, the guest kernel runs in Ring 1, and applications, as usual, run in Ring 3. In 64-bit mode, Xen runs in Ring 0, while the guest kernel and applications run in Ring 3 due to lack of segmentation support.

Virtualising the MMU: As x86-based CPUs do hardware page table walk to fill in TLB miss, it is hard to virtualise MMU without losing out on performance. That is where paravirtualisation helps a lot. Here, guests have direct access to the physical pages and they continue to build and maintain the page tables. Xen steps in to ensure that page table pages don't get modified without its knowledge, so it marks them as 'read only'. Any updates to them by the guest kernel are trapped and validated before being updated.

Virtualising the IDT: The `set_trap_table()` hypercall is used by the guests to register their exception handlers. Whenever an exception happens, a stub routine in the Xen hypervisor executes. This will pass on the exception to the appropriate guest via event channels. Then the corresponding routine from the above trap table is invoked. Two major faults that happen very frequently on a system are: system calls and pagefaults. The throughput of the system is dependent on how these two faults are handled. Xen installs guest-specific handlers directly in the IDT, thus reducing one level of indirection in the fault handler. Xen validates these handlers before installing them in the IDT.

Virtualising the IO: This is another area where paravirtualisation helps in boosting the performance of the guest domains. Guest domains do know that they don't have direct control of the hardware. Xen provides a '*split driver model*', which gives close to native performance in IO. This works by providing a class-specific *frontend driver* in each of the guest domains; it talks to a '*backend driver*' available in dom0. [Xen also provides a mechanism whereby a particular domU can also host a peripheral—in which case that domU will provide the back-end driver.] The front-end and back-end drivers communicate using *event channels*.

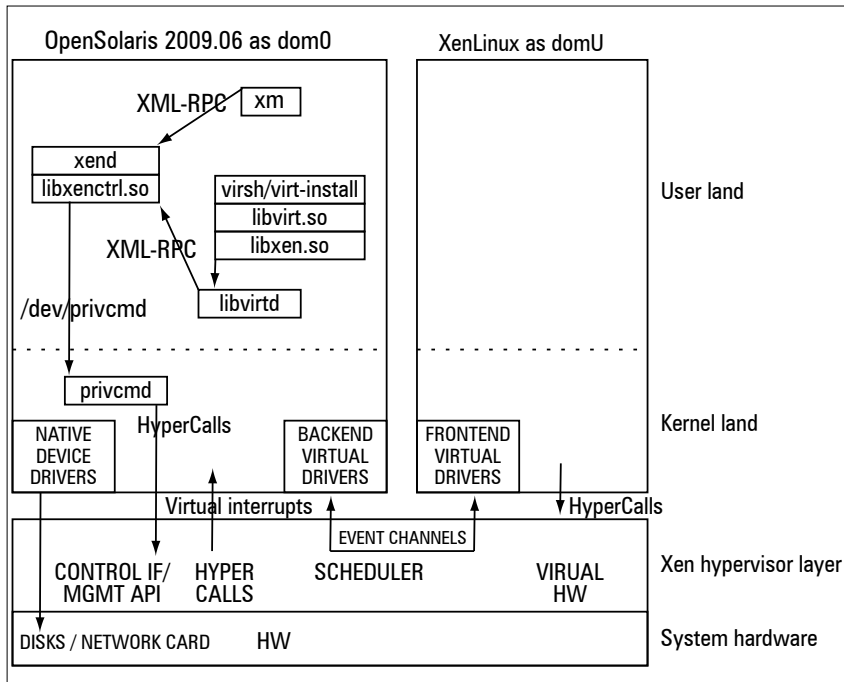


Figure 1: Thin hypervisor sitting on the hardware

The back-end driver, in turn, communicates with the actual driver that controls the device. What that means is, generally, all the peripheral devices are controlled by drivers in dom0, rather than being present in the hypervisor itself. This is a very elegant idea to keep the hypervisor slick and fast. Front-end drivers are class-specific—i.e., there will be one FE for block devices, irrespective of whether the actual device is an IDE, an SCSI or FC. Likewise, there will be only one FE for network devices.

Virtual interrupts: Just as hypercalls provide guest-to-Xen communication, 'event channels' provide Xen-to-guest communication. They are the software equivalent of hardware interrupts. They provide an asynchronous notification mechanism. An event channel is associated with every virtual device. IO request/completion notifications happen via event channels.

Start info page: A PV guest booting in the Xen environment does not have access to the firmware. This is where the 'start info' page kicks in. It's a page mapped into the guest address space by the hypervisor with

details on how much memory the guest is configured with, XenStore details, console details and another page called the 'shared info' page. The `/usr/include/xen/xen.h` file contains the declaration of 'struct start_info'. A guest kernel starts off with information present in this page and then it maps in the XenStore page to get all the devices configured for it. It also gets the *console page* to print out any debug information and thus troubleshoots any bugs present early in the boot phase. A PV guest starts in protected mode unlike a normal OS, which starts in real mode.

Shared info page: Any OS starts off with a single CPU and gets the other CPUs online in an SMP system later in the boot process. Likewise, even a PV guest starts off with a single CPU and the rest of the CPU information is retrieved from the 'shared info' page. It also contains architecture-specific information, if any. Besides this, the 'shared info' page also contains asynchronous notifications pending to the guest as well as those masked by the guest—think of them as signals


pending to a process and those masked by a process. A guest has to explicitly map in the 'shared info' page before it could access the contents.

Balloon thread: This runs in a guest kernel and coordinates with the hypervisor to get more physical memory to a guest; or in the case of overall system pressure, gives back some pages to the hypervisor.

Scheduling: The Xen hypervisor uses the *credit scheduler* to multi-task virtual CPUs onto the physical CPUs. The *credit scheduler* is more like the *fair share scheduler* you will find on OpenSolaris.

Time keeping: When a guest goes off the CPU and comes back later, it has no clue how much *wall clock time* lapsed in between. To handle this, hypervisor stores the *walk clock time* when the guest actually booted in the *shared info page*. The *virtual clock time* of a guest increments while the guest is active. Using these two, the current *wall clock time* is calculated. A PV guest can also sync with a NTP source to keep *wall clock time* accurate.

The discussion thus far is encapsulated in Figure 1. Here, two guests are running simultaneously on the system. One is the privileged guest dom0, which hosts quite a few services for the hypervisor, provides drivers for the peripherals, and hosts the back-end drivers. The second is an unprivileged guest that talks to the devices via the front-end drivers.

Hopefully, we are now clear on many of the essential concepts associated with Xen virtualisation technology. In Part 2 of this series (turn to page 69), we move on to setting up OpenSolaris 2009.06 as dom0. 

By: Surya Prakkhi

Surya debugs OpenSolaris kernel, both for living and for passion, and relies a lot on the lifeline called `dtrace(1M)`. He occasionally blogs at blogs.sun.com/sprakkhi.

OpenSolaris on Xen **Part 2**

Exploring a

dom0

OpenSolaris

Let's play around with a dom0 OpenSolaris to understand it better.

In this article we'll use the latest version of OpenSolaris (version 2009.06) to set up Xen. Refer to '*Solaris For Your Desktop*', published in the April 2009 issue of LFY, if you need installation instructions. Additionally, since version 2009.06 doesn't come with the Xen packages installed and we need these to be able to boot the system as Xen dom0, follow the instructions given at http://opensolaris.org/os/community/xen/docs/2008_11_dom0 to achieve the same.

Upon reboot, pick up the '*xvm*' Grub entry to boot into Xen. Once it comes up, log in as the user you configured while installing os2009.06. To make sure that this guest is running on the Xen hypervisor, follow the steps below:

```
surya@opensolaris:~$ uname -i
i86xpv
surya@opensolaris:~$
```

If *uname -i* reports i86xpv, then it is a PV guest. What does it report on bare metal? To confirm that this is a privileged domain (i.e., dom0), run the following:

```
surya@opensolaris:~$ pfexec cat /dev/xen/domcaps
control_d
surya@opensolaris:~$
```

This confirms it is a '*control domain*'. For any other domU, it will be empty.

To get a feel of how 'thin' the hypervisor is, look at the hypercalls listed in */usr/include/xen/xen.h*—they add up to less than

40. To put this number in perspective, OpenSolaris provides more than 200 system calls for applications, to access kernel services.

SMF @ work

If the following six SMF services are online, that means we are ready to install new guests on this Xen system. You can run `svcs -p` against each of those services to see what processes are associated with those services (for example, `#svcs -p xvm/store`). As these long-running daemons are wrapped as SMF services, benefits like an automatic restart in case of unexpected termination, follow.

```
surya@opensolaris:~$ svcs -a | grep xvm
disabled Jun_26 svc:/system/xvm/ipagent:default
online   Jun_26 svc:/system/xvm/vnc-config:default
online   Jun_26 svc:/system/xvm/store:default
online   Jun_26 svc:/system/xvm/xend:default
online   Jun_26 svc:/system/xvm/console:default
online   Jun_26 svc:/system/xvm/virtld:default
online   Jun_26 svc:/system/xvm/domains:default
surya@opensolaris:~$
```

Note that `'ipagent'` doesn't run in dom0; it runs only in PV domUs.

The power of ZFS

We could create a ZFS volume and pass it as an install disk to a guest and reap the benefits associated with this modern filesystem, such as cloning (we can deploy many more VMs after modifying per instance data), snapshots, end-to-end data integrity, unlimited space, etc.

We use `virt-install` (1) to install new guests. Please follow the instructions at http://opensolaris.org/os/community/xen/docs/opensolaris_domu to install OpenSolaris 2009.06 as a PV guest. After a successful install and reboot of the new domU, from dom0 we can list out the currently installed domains, using the following:

```
surya@opensolaris:~$ pfexec virsh list
Id Name      State
-----
0 Domain-0   running
1 osol       blocked
surya@opensolaris:~$
```

As you can see, I have named my OpenSolaris 2009.06 guest as 'osol'. Now you can take the domU console using `'virsh console'`. Log in to the domU using the account we configured while installing. Log in and confirm that it is not a privileged domain.

Thus we now have a domU running besides dom0—i.e., two operating systems running simultaneously on a

single physical machine—Xen at work!

Dynamically, we can add more disk capacity to the osol domU by passing additional files or block devices from dom0.

First, confirm from inside domU, how many block devices are available using the `format (1)` command. We can use either `'virsh (1M)'` or `'xm (1M)'` to attach new drives from dom0. `xm (1M)` stands for Xen Manager and is the original tool that came along with Xen. It provides command and control of the hypervisor. This is being deprecated in favour of `virsh`, albeit in places where `virsh` hasn't yet provided the same functionality.

```
surya@opensolaris:~$ mkfile 100m /var/tmp/disk2
surya@opensolaris:~$ pfexec chown xvm /var/tmp/disk2
surya@opensolaris:~$ pfexec xm block-attach osol tap:vdisk:/var/tmp/
disk2 3 w
surya@opensolaris:~$
surya@opensolaris:~$ pfexec zfs create -V 5G rpool/xvm-drive3
surya@opensolaris:~$ pfexec virsh attach-disk osol /dev/zvol/dsk/rpool/
xvm-drive3 xvdd --driver phy
surya@opensolaris:~$
```

The above commands add a regular file and ZFS volume to the domU as additional disks.

From dom0, we could use `'virsh dumpxml osol'` and spot the new disks that we just added. From domU, we could run the `format (1M)` command and spot the presence of additional drives. We could run `pfexec virsh detach-disk xvdd` and see one of the drives vanish from domU.

As can be seen, the changes are reflected in the domU immediately.

Similarly, additional vNICs can also be passed onto the running guest. Likewise, vCPUs can also be taken off the running domain. But to pass more vCPUs than what is already configured, domU has to be shut down and its config file updated—i.e., run `'virsh dumpxml osol > /tmp/osol.xml'` and try changing vCPUs in the `/tmp/osol.xml` file, then undefine and redefine the guest using the new `/tmp/osol.xml` file.

Secure by default

Look for a process called `'libvirtld'` in dom0 and `truss (1)` it:

```
surya@opensolaris:~$ pfexec truss -l -o /tmp/foo -p `pgrep libvirtld`
```

Now, from another dom0 terminal, execute `pfexec virsh list`, and the above `truss` starts tracing all the system calls being made by `libvirtld`. So what is happening? All the `libvirt` tools communicate with `xend` via `libvirtld`—the `libvirtld` process is made to run with *least privileges* (5), whereas `virsh (1M)` and family are not.

We could run `ps -f -u xvm` in dom0 and confirm that none of the services associated with Xen run

as the root. They run as ordinary user *xvm* with an appropriate privilege. This way, OpenSolaris complies with security out-of-the-box.

DTrace

We can trace the IO path *read* (2) takes in a domU using *dtrace* (1M). Knowing that the entry point into a block driver is the strategy routine, let us go after *xdf_strategy()*, the strategy routine of the block front-end driver:

```
surya@opensolaris:~$ pfexec dtrace -n xdf::entry' {@[stack()]=count()}'
dtrace: description 'xdf_strategy:entry' matched 1 probe
```

Meanwhile, from a domU terminal, *cat* (1) a text file; then come back to the DTrace terminal and Ctrl+C it; you will find a stack similar to the one below:

```
genunix`bdev_strategy+0x75
genunix`ldi_strategy+0x59
zfs`vdev_disk_io_start+0xd3
zfs`zio_vdev_io_start+0x17d
zfs`zio_execute+0xa0
zfs`zio_nowait+0x5a
zfs`vdev_mirror_io_start+0x148
zfs`zio_vdev_io_start+0x1ba
zfs`zio_execute+0xa0
zfs`zio_nowait+0x5a
zfs`arc_read_nolock+0x82a
zfs`arc_read+0x75
zfs`dbuf_read_impl+0x159
zfs`dbuf_read+0xfe
zfs`dmu_buf_hold_array_by_dnode+0x1bf
zfs`dmu_buf_hold_array+0x73
zfs`dmu_read_uio+0x4d
zfs`zfs_read+0x19a
genunix`fop_read+0x6b
genunix`read+0x2b8
1
```

This stack gives you enough details on how the read call is directed through different ZFS routines before *xdf* is invoked. Knowing that, ZFS calls the *zfs_interrupt* routine whenever an IO finishes. Let's trace that call too—we'll get a stack trace similar to the following:

```
zio_interrupt:entry
zfs`vdev_disk_io_intr+0x6b
genunix`biodone+0x84
xdf`xdf_io_fini+0x93
xdf`xdf_intr_locked+0x78
xdf`xdf_intr+0x2b
unix`av_dispatch_autovect+0x7c
unix`dispatch_hardint+0x37
unix`switch_sp_and_call+0x13
```

This gives us an understanding of what happens when the notification comes in upon IO completion.

Similarly, we can trace the IO path in dom0 as well, to see how the Xen back-end driver talks to the actual driver by going after the *xdb_biodone* call. Likewise, we can observe the network IO paths too—albeit by tracing STREAMS routines.

There is a new DTrace provider, *xdt*, which monitors hypervisor events like vCPU switching. All the probes associated with it can be found by using *dtrace -l -P xdt*.

After observing o/p from the above DTrace command, we can understand why the status of a domU, as reported by *virsh list*, is blocked—because the vCPU of *osol* domU is not currently running on any physical CPU.

We can use */usr/lib/xen/bin/xenstore-snoop* to snoop on the interaction between dom0 and domU via XenStore. We can run this command in a dom0 terminal, and from another dom0 terminal just run *virsh dumpxml osol* and see what *xenstore-snoop* reports.


Similarly, run the *format* command from a domU terminal and see what the above snoop reports. For the curious, this command is implemented using the DTrace PID provider. This command comes in handy to keep an eye on XenStore if we perceive any slowness in the system.

Crash file support

In case the hypervisor panics, OpenSolaris dom0 is extended to collect the core of the whole system in which Xen itself will be available as a kernel module. More information on this is available at http://blogs.sun.com/nilsn/entry/debugging_an_xvm_panic.

Secondly, if we find that a PV OpenSolaris guest is hung (with no response on console and not reachable over the network), we could collect the core of the domU from dom0 as follows, and use *mdb* (1) on the core file to investigate further:

```
surya@opensolaris:~$ pfexec xm dump-core osol
Dumping core of domain: osol ...
surya@opensolaris:~$ pfexec mdb /var/xen/dump/2009-0701-1528.42-
osol.3.core
Loading modules: [ unix genunix specs dtrace mac .. logindmux ufs
nsmdb sppp mpt emlxs ]
>
```

We have just scratched the surface of the magic that can be done with OS2009.06 as dom0. Explore more at www.opensolaris.org/os. 

By: Surya Prakk

Surya debugs OpenSolaris kernel, both for living and for passion, and relies a lot on the lifeline called *dtrace*(1M). He occasionally blogs at blogs.sun.com/sprakki.



The New Scheduler

on the Block, Dedicated to Desktops

What in the world is the BFS and how does it fare against the mainline kernel's CFS? Here's a report.

After a two-year break from kernel hacking (during which time he probably was hacking Linux in private), Australian anaesthetist and part-time kernel developer, Con Kolivas, has emerged with a new scheduler. He claims it provides significantly better performance on dual- and quad-core processor-based desktops compared to the mainline kernel's Completely Fair Scheduler (CFS), developed by Ingo Molnar. (Read more about Kolivas' past involvement with kernel development and why he had quit in 2007 at www.linuxforu.com/views/sad-case-of-dr-con.)

Kolivas has named it the Brain Fuck Scheduler, or BFS, for short.

Why the name?

There are numerous reasons for this. According to the official FAQ, here are a few:

- Because it throws out everything that we know is good about how to design a modern scheduler in scalability.
- Because it's so ridiculously simple.
- Because it performs so ridiculously well in what it's good at, despite being that simple.
- Because it's designed in such a way that

mainline would never be interested in adopting it, which is how I like it.

- Because it will make people sit up and take notice of where the problems are in the current design.
- Because it throws out the philosophy that one scheduler fits all and shows that you can do a *lot* better with a scheduler designed for a particular purpose. I don't want to use a steamroller to crack nuts.
- Because it actually means that more CPUs mean better latencies.
- Because I must be fucked in the head to be working on this again.
- I'll think of some more because later.

Now, before we get deep into BFS and its performance benefits on desktops, let's have some basic ideas about schedulers.

What's a scheduler?

The sole purpose of an operating system is to multi-task. The number of tasks (i.e., processes running in the foreground, as well as in the background) are always far greater than the number of CPUs in our computers. How does the CPU(s) attend to processing requests, which often come simultaneously, from a number of processes?

This is where a software called ‘scheduler’ comes into the picture—to schedule running processes to CPU(s) time. Thus, it is the scheduler that makes it possible for programs to run concurrently in an operating system.

There are several terms associated with a process scheduler—some of them, to quote Wikipedia, are:

- CPU utilisation: To keep the CPU as busy as possible.
- Throughput: The number of processes that complete their execution per time unit.
- Turnaround: The total time between submission of a process and its completion.
- Waiting time: The amount of time a process has been waiting in the ready queue.
- Response time: The amount of time it takes from when a request was submitted until the first response is produced.
- Fairness: Equal CPU time to each thread.

Jargon apart, the purpose of a scheduler is to always provide a responsive user experience by allocating tasks to the CPU efficiently.

Linux schedulers

Linux has had a multi-level feedback queue with priority levels ranging from 0-140 since version 2.5. Priority levels in the range 0-99 are reserved for real-time tasks and 100-140 are considered nice task levels. From versions 2.6.0 to 2.6.23, Linux used an $O(1)$ scheduler that could schedule processes within a constant amount of time, regardless of how many processes are running on the operating system. When kernel version 2.6.23 was released, developers replaced this scheduler with the Completely Fair Scheduler that uses Red Black trees instead of queues.

Completely Fair Scheduler

Ingo Molnar developed the Completely Fair Scheduler, inspired by Kolivas' implementation of fair scheduling called Rotating Staircase Deadline. Quoting Wikipedia, “CFS is an implementation of a well studied, classic algorithm called fair queuing. CFS does not consider a task to be a sleeper if it sleeps for very short time—a short sleeper might be entitled to some bonus time, but never more than it would have had, had it not slept. CFS has a scheduling complexity of $O(\log N)$ where N is the number of tasks in the run queue.”

Run queues in CFS are implemented as Red Black Trees [en.wikipedia.org/wiki/Red-black_tree]. Because of this implementation, the selection of a process for CPU allocation can be done in constant time, but reinserting the job after it has run needs $O(\log N)$ operations.

The name ‘Completely Fair Scheduler’ is technically incorrect as the algorithm only guarantees the ‘unfair’ level to be less than $O(n)$ where n is the number of processes. There are more complicated algorithms that can perform well at ‘unfair’ levels—for example, $O(\log n)$. CFS helps assure every process to get its fair share of CPU time.

Levels of scheduling

In multi-tasking systems, schedulers operate at three levels based on the relative frequency with which they perform their functions. The different levels of scheduling are:

- **Long-term scheduling:** The long-term, or admission, scheduler decides which programs are admitted to the system for execution and when, and which ones should be exited. In modern operating systems, this is used to make sure that real-time processes get enough CPU time to finish their tasks. Long-term scheduling is also important in large-scale systems such as batch processing systems, computer clusters, super computers and render farms.
- **Mid-term scheduling:** The mid-term scheduler makes use of virtual memory, by temporarily removing processes from main memory and places them on secondary memory (such as a disk drive) or vice versa. This technique is known as “swapping out” or “swapping in”. The mid-term scheduler decides which processes are needed to be swapped by finding out the processes which have not been active for some time, or a process which has a low priority, or a process which is page faulting frequently, or a process which is taking up a large amount of memory in order to free up main memory for other processes, swapping the process back in later when more memory is available, or when the process has been unblocked and is no longer waiting for a resource.
- **Short-term scheduling:** The short-term scheduler or dispatcher decides which of the ready, in-memory processes are to be executed next and for how long. In this way, the short-term scheduler makes scheduling decisions much more frequently than the long-term or mid-term schedulers. This scheduler can be preemptive, which means that it can forcibly remove processes from a CPU when it wants to allocate that CPU to another process, or non-preemptive—in which case the scheduler is unable to forcibly deallocate the processes off the CPU.

—Courtesy: Wikipedia [[en.wikipedia.org/wiki/Scheduling_\(computing\)](http://en.wikipedia.org/wiki/Scheduling_(computing))]

What's with this BFS?

The xkcd.com comic ‘Supported_Features’ [Figure 1; <http://xkcd.com/619/>], which picked on Linux for having support for 4,096 CPUs while not supporting essential desktop computing features like smooth flicker-free full screen Flash videos, inspired Kolivas to start working on a new scheduler to address the requirements of desktop users.

Kolivas wrote BFS to improve responsiveness on non-uniform memory access (NUMA) Linux desktop computers



Figure 1: xkcd.com comic picks on Linux

and mobile devices with less than 16 cores. According to him, the current kernel scheduler, CFS, fails to get the maximum out of the hardware in most common situations on a desktop.

As Kolivas writes in his FAQ page for BFS: "For years we've been doing our workloads on Linux to have more work than we had CPUs because we thought that the 'jobserver' were limited in their ability to utilise the CPUs effectively (so we did make `-j6` or more on a quad core machine, for example). This scheduler proves that the job servers weren't at fault at all, because make `-j4` on a quad core machine with BFS is faster than *any* choice of job numbers on CFS." The difference in speed is attributed to the inability of the mainline scheduler to keep the CPUs busy enough like BFS does.

Highlighting the design goals of BFS, he writes, "It was designed to be forward looking only, to make the most of lower spec machines, and not scale to massive hardware. That is, it is a desktop-orientated scheduler, with extremely low latencies for excellent interactivity by design rather than 'calculated', with rigid fairness, nice priority distribution and extreme scalability within normal load levels."

Thus BFS strives to produce extreme low latencies on machines with up to 16 CPUs. Anything more than this many CPUs will have exponentially less performance. So BFS may not perform well, or even perform worse, on huge machines with more number of CPUs. As the xkcd comic rightfully puts it: "But, who uses that?"

So, what other striking features does BFS have? Kolivas writes, "On top of the current scheduler design, it has a `SCHED_IDLEPRIO` which actually does only schedule tasks when idle, and `SCHED_ISO` for unprivileged real-time performance. BFS does NOT implement `CGROUPS`. A desktop user should not need to know about `CGROUPS`, nor should they need to use them. BFS also does not have the feature of 'lots of tunables I don't understand'."

By removing most of the unnecessary features that a desktop computer doesn't require, it is the simplicity that works in favour of BFS.

How do I use BFS?

You may need to use one of the latest versions of the kernel and patch it with the appropriate BFS patch. Go and fetch the latest kernel (I choose version 2.6.31) from kernel.org:

```
cd /usr/src
```

```
wget -c http://www.kernel.org/pub/linux/kernel/v2.6/linux-2.6.31.tar.bz2
```

Then download the BFS patch for that kernel and apply the patch after extracting the kernel source:

```
wget -c http://ck.kolivas.org/patches/bfs/2.6.31-sched-bfs-220.patch
```

```
tar xjvf linux-2.6.31.tar.bz2
```

```
cd linux-2.6.31
```

```
patch -p1 < ../2.6.31-sched-bfs-220.patch
```

Then compile the kernel. If you are using Ubuntu, you can use the `make-kpkg` command to build the deb packages for *kernel-image* and *kernel-headers*. You may need to install *kernel-package* prior to this. It is as simple as typing a `sudo apt-get install kernel-package`. You can copy the current kernel configuration as follows:

```
cp -vi /boot/config-`uname -r` .config
```

Note that since Ubuntu kernels are built with *Debugging info* enabled, it makes the kernel modules much larger in size than they should be. So, if you're using the existing configuration, it's advisable to turn off the *Debugging info* mode. While you're at `make menuconfig`, go to 'Kernel hacking', and under 'Kernel debugging' turn off 'Compile the kernel with debug info'.

If you want to speed up the compilation process, you can enable parallel make by setting the `CONCURRENCY_LEVEL` environment variable. Set it to 1 plus the number of processor cores you have. For example, if you have a dual-core processor, use:

```
export CONCURRENCY_LEVEL=3
```

Change the kernel configuration:

```
make menuconfig
```

Additionally, if you're not compiling for the first time, run the following command for a clean build:

```
make-kpkg clean
```

Finally:

```
fakeroot make-kpkg --initrd --append-to-version=-bfskernel kernel-image kernel-headers
```

This will take a long time and after this step, you will get two files, namely:

- `linux-headers-2.6.31-bfs211-bfskernel_2.6.31-bfs211-bfskernel-10.00.Custom_i386.deb`
- `linux-image-2.6.31-bfs211-bfskernel_2.6.31-bfs211-bfskernel-10.00.Custom_i386.deb`

Install them using the `dpkg -i <deb file>` command.




Figure 2: The reverse scalability graph shows various job numbers on a kernel build on a quad core machine

Then reboot your machine to the new kernel. You will get a new kernel entry in Grub that looks like "Ubuntu 9.04, kernel 2.6.31-bfs211-bfskernel". Log in and experience the difference.

Performance analysis

As Kolivas explains in his FAQ page, it is not the number of physical CPUs that affect BFS's scalability. Rather, it's that of logical ones. This means that a hyper-threaded quad core CPU is simulated as 8 logical CPUs.

People are also concerned about the $O(n)$ lookup of BFS. However, the actual overhead is very small. Even with large numbers of n , it has lower overhead than an $O(1)$ design. During a test against CFS, in which the test application forks 1,000 simple tasks, it showed no difference in time as compared to CFS on a quad core machine. But BFS will be much faster when the load is lower and the CPUs are less in number, which is normally the case with desktops.

Kolivas cites an example in the FAQ page, "See the reverse scalability graph, courtesy of Serge Belyshev, showing various job numbers on a kernel build on a quad core machine [Figure 2]. The problem has always been that the mainline scheduler can't keep the CPUs busy enough; i.e., it doesn't make the most of your hardware in the most common situations on a desktop! Note that the reverse scalability graph is old; the scalability has improved since then."

To test the speed of the new scheduler, I compiled Python 2.6.2 with a `make -j2` on my laptop that has a Core2Duo processor and 2GB of RAM. It took only 1 minute and 34 seconds to complete on the new BFS patched kernel. The compilation using simple `make` (without the `-jX` option) completed in 2 minutes and 2 seconds.

Earlier, I had done this test on the stock kernel that uses CFS. But it took 1 minute and 41 seconds using `make -j2` and 2 minutes and 4 seconds without using the `-jX` option to compile the same version of Python. So the speed difference is obvious. Refer to Tables 1 and 2 for the time differences.


make -j2 performance while compiling Python 2.6.2		
	when using BFS	when using CFS
real	1 minute, 34.297 seconds	1 minute 41.841 seconds
user	1 minute, 56.151 seconds	1 minutes 57.387 seconds
sys	7.836 seconds	8.309 seconds

Table 1

make performance while compiling Python 2.6.2		
	when using BFS	when using CFS
real	2 minutes, 2.595 seconds	2 minutes, 4.598 seconds
user	1 minute, 54.183 seconds	1 minute, 54.519 seconds
sys	7.576 seconds	7.836 seconds

Table 2

Some users from the Ubuntu Forums reported that they have experienced a speed difference in GDM. "Before the patch, there was a period of some five seconds during which GDM seemed to be in deep thought mode, trying to decide which users to show in this list. After the patch the list appeared almost immediately."

There have been only positive reviews of the BFS scheduler throughout the Linux community, as far as I could see. As the developer says, BFS is not stable yet and when it reaches a stable release, I'm pretty sure it will have something more to offer. By the way, Kolivas concludes his FAQ page by pointing out, "Yes I know you probably still can't watch full screen videos on YouTube, but that's not entirely the scheduler's fault." 

References

- Wikipedia on schedulers: [en.wikipedia.org/wiki/Scheduling_\(computing\)](http://en.wikipedia.org/wiki/Scheduling_(computing))
- Types of scheduling: <http://www.cim.mcgill.ca/~franco/OpSys-304-427/lecture-notes/node38.html>
- Wikipedia on CFS: en.wikipedia.org/wiki/Completely_Fair_Scheduler
- BFS FAQ: ck.kolivas.org/patches/bfs/bfs-faq.txt
- Con Kolivas interview: ck.kolivas.org/german_linux_magazine_interview.txt
- The reverse scalability graph is taken from ck.kolivas.org/patches/bfs/reverse-scalability.png

By: Sreejith K.

The author is a FOSS enthusiast and a hardcore Ubuntu fan who loves hacking GNU/Linux. He's currently working as a development engineer at K7 Computing, Chennai. His areas of interest include compilers, virtualisation technologies and cloud computing. He blogs at <http://semk.in> and can be reached at [sreejithemk AT gmail DOT com](mailto:sreejithemk@gmail.com).

Creating VPN with OpenVPN

In this article we'll take a look at two of the most commonly used configurations for OpenVPN.



Virtual Private Networks, or VPN for short, allow us to establish a private network over a public network (like the Internet) using cryptography to secure it. An SSL VPN uses an SSL (Secure Sockets Layer), enabling the service to be accessible over the HTTPS protocol using a Web browser. OpenVPN is an example of an SSL VPN solution that's open source and is targeted at SMEs and enterprises.

A wide range of configurations exist for OpenVPN, including site-to-site VPNs,

remote access, Wi-Fi security, etc. It uses the TUN interface for virtual linking of the two end points of the VPN. The TUN interface looks like a point-to-point network hardware device to the operating system, but is, in fact, nothing but a virtual network adapter. Thus being a virtual interface, the data is actually pushed to the user-space program, which can then open a TUN device (like a file), and read and write IP packets from and to it.

Note that a TAP interface is similar to TUN except that it emulates the Ethernet

rather than a point-to-point network.

In a user-space VPN, the IP packets from a TUN/TAP adapter are encrypted and then encapsulated into UDP and sent over the Internet. At the destination, the remote host de-encapsulates, decrypts and authenticates these IP packets before pushing them into a TUN/TAP virtual adapter.

Configuring OpenVPN

There are various ways to configure OpenVPN and you can choose the one that matches your requirement. In this article we will look at a few types of configurations that are typically used.

To start with, install OpenVPN as follows. The server set-up in this article uses a Linux system, while various clients can be used for testing purposes, like Windows XP with `openvpn-gui`; Fedora and Ubuntu with an OpenVPN plug-in for NetworkManager; as well as Debian, CentOS and Red Hat with command line-based OpenVPN clients.

We will start with installing the software from the distributions repository.

On CentOS 5.3 I had activated the Fedora EPEL repository [<https://fedoraproject.org/wiki/EPEL>] for the OpenVPN package to make sure that I used a 2.1 version build.

```
# yum install openvpn
```

The configuration files are kept in the `/etc/openvpn` directory. The `init` script installed with the package takes care of registering the OpenVPN with the `chkconfig` system, so that it can be started and stopped using the `service` command.

While starting the `openvpn` service, the `init` script takes care of loading all the configuration files from the `/etc/openvpn` directory with the `.conf` extension. If a particular configuration requires some initialisation to be done, then we can create a script in `/etc/openvpn/` by the same name as the configuration file with a `.sh` extension. So, for example, let's call our configuration file `/etc/openvpn/server.conf`. We can create a script `/etc/openvpn/server.sh` that will be executed before loading the configuration from the `server.conf` file.

We will use the following directory structure for our set-up, so let us create it:

```
# mkdir /etc/openvpn/keys
```

Regardless of the configuration chosen, there are a few options used in all types of configurations. Although some of these are default settings, it is a good idea to specify them in the configuration file. The following example is the static key configuration, which is the most basic type, so I will describe everything about it thoroughly; from thereon, for other configurations, just the additions and removals will be specified.

Initially we only allow the client and server to ping each other. After reviewing all popular configurations we will have a look at how to provide the clients access to the private network behind the VPN, and also how to configure a proxy to let the clients surf the Net securely.



Note: On the CentOS Server SELinux was enabled, so the server set-up described here should work on the SELinux-enabled boxes. At certain places I have shown the output displaying the SELinux contexts.

Static key configuration

The static key configuration of OpenVPN is the most basic type and only allows one client to connect to one server. You can use this simple set-up if you want to set up VPN connectivity between your laptop or home computer and one of your servers on the Internet, somewhere. This is the quickest configuration to set up. However, we need to make sure that after the key is generated at one end, it is securely transferred to the other end before initiating the connection. This is only a one-time job, so the effort is worth it.

To start with, we will generate the static key:

```
root@vpn.unixclinic.net # cd /etc/openvpn/keys
root@vpn.unixclinic.net # /usr/sbin/openvpn --genkey --secret static.key
root@vpn.unixclinic.net # ls -lZ
-rw----- root root user_u:object_r:openvpn_etc_rw_t static.key
```

Now let us configure the server part:

```
root@vpn.unixclinic.net # vi /etc/openvpn/static-server.conf
port 1194
proto udp
dev tun

# The keep alive directive is particularly important if you are using UDP
# through a stateful firewall like Netfilter. Because UDP is connectionless
# any stateful firewall will forget about the connection if packets are not
# going through it at regular intervals.
keepalive 10 60

ping-timer-rem
persist-tun
persist-key

# Enable compression (use only if compiled with lzo support)
comp-lzo

# Short log of active connections and internal routing table.
# Recreated every minute.
status openvpn-status.log

log-append openvpn.log
```

```
# Verbosity level in the log (0=silent, 3 or 4=normal, 9=maximum for
debugging)
verb 3
### Change following for different configurations
# For Static key server configuration
ifconfig 10.8.0.1 10.8.0.2
secret /etc/openvpn/keys/static.key
```

At the client side, the configuration file will look like the following code:

```
root@client # vi /etc/openvpn/static-client.conf
port 1194
proto udp
dev tun

# The keep alive directive is particularly important if you are using UDP
# through a stateful firewall like Netfilter. Because UDP is connectionless
# any stateful firewall will forget about the connection if packets are not
# going through it at regular intervals.
keepalive 10 60

ping-timer-rem
persist-tun
persist-key

# Enable compression (use only if compiled with lzo support)
comp-lzo

# Short log of active connections and internal routing table.
# Recreated every minute.
status openvpn-status.log

log-append openvpn.log

# Verbosity level in the log (0=silent, 3 or 4=normal, 9=maximum for
debugging)
verb 3

# Remote OpenVPN server to connect to
remote vpn.unixclinic.net

### Change following for different configurations
# For Static key client configuration
ifconfig 10.8.0.2 10.8.0.1
secret /etc/openvpn/keys/static.key
```

As you can see, there is only one line that is very different, which is the *ifconfig* line, where the IP addresses are reversed. Also, there is a configuration option named 'remote' at the client side. The purpose of this option is to specify the OpenVPN server to which this client has to connect with. The 'remote' option can take both the IP address and the resolvable hostname. Please note that you need to find a pre-existing secure

channel for the static key to be transferred to the client from the server.

Now start the OpenVPN server at both the ends, and *tail* the log files at both ends for any messages:

```
root@vpn.unixclinic.net # tail -50 /etc/openvpn/openvpn.log
Thu Jun 18 03:16:04 2009 OpenVPN 2.1_rc15 x86_64-redhat-linux-gnu
[SSL] [LZO2] [EPOLL] built on Nov 30 2008
Thu Jun 18 03:16:04 2009 NOTE: the current --script-security setting may
allow this configuration to call user-defined scripts
Thu Jun 18 03:16:04 2009 Static Encrypt: Cipher 'BF-CBC' initialized with
128 bit key
Thu Jun 18 03:16:04 2009 Static Encrypt: Using 160 bit message hash
'SHA1' for HMAC authentication
Thu Jun 18 03:16:04 2009 Static Decrypt: Cipher 'BF-CBC' initialized with
128 bit key
Thu Jun 18 03:16:04 2009 Static Decrypt: Using 160 bit message hash
'SHA1' for HMAC authentication
Thu Jun 18 03:16:04 2009 LZO compression initialized
Thu Jun 18 03:16:04 2009 TUN/TAP device tun0 opened
Thu Jun 18 03:16:04 2009 TUN/TAP TX queue length set to 100
Thu Jun 18 03:16:04 2009 /sbin/ip link set dev tun0 up mtu 1500
Thu Jun 18 03:16:04 2009 /sbin/ip addr add dev tun0 local 10.8.0.1 peer
10.8.0.2
Thu Jun 18 03:16:04 2009 Data Channel MTU parms [ L:1545 D:1450 EF:45
EB:135 ET:0 EL:0 AF:3/1 ]
Thu Jun 18 03:16:04 2009 Local Options hash (VER=V4): '4b91e501'
Thu Jun 18 03:16:04 2009 Expected Remote Options hash (VER=V4):
'48593abd'
Thu Jun 18 03:16:04 2009 Socket Buffers: R=[124928->131072] S=[124928-
>131072]
Thu Jun 18 03:16:04 2009 UDPv4 link local (bound): [undef]:1194
Thu Jun 18 03:16:04 2009 UDPv4 link remote: [undef]
Thu Jun 18 03:16:09 2009 Peer Connection Initiated with
116.87.186.181:1194
Thu Jun 18 03:16:10 2009 Initialization Sequence Completed
Thu Jun 18 03:18:46 2009 write UDPv4 [EHOSTUNREACH]: Operation not
permitted (code=1)
```

```
root@client # tail -50 /etc/openvpn/openvpn.log
Thu Jun 18 15:16:08 2009 OpenVPN 2.1_rc15 x86_64-redhat-linux-gnu
[SSL] [LZO2] [EPOLL] built on Nov 30 2008
Thu Jun 18 15:16:08 2009 NOTE: the current --script-security setting may
allow this configuration to call user-defined scripts
Thu Jun 18 15:16:08 2009 Static Encrypt: Cipher 'BF-CBC' initialized with
128 bit key
Thu Jun 18 15:16:08 2009 Static Encrypt: Using 160 bit message hash
'SHA1' for HMAC authentication
Thu Jun 18 15:16:08 2009 Static Decrypt: Cipher 'BF-CBC' initialized with
128 bit key
Thu Jun 18 15:16:08 2009 Static Decrypt: Using 160 bit message hash
'SHA1' for HMAC authentication
Thu Jun 18 15:16:08 2009 LZO compression initialized
Thu Jun 18 15:16:08 2009 TUN/TAP device tun0 opened
Thu Jun 18 15:16:08 2009 TUN/TAP TX queue length set to 100
```



```
Thu Jun 18 15:16:08 2009 /sbin/ip link set dev tun0 up mtu 1500
Thu Jun 18 15:16:08 2009 /sbin/ip addr add dev tun0 local 10.8.0.2 peer
10.8.0.1
Thu Jun 18 15:16:08 2009 Data Channel MTU parms [ L:1545 D:1450 EF:45
EB:135 ET:0 EL:0 AF:3/1 ]
Thu Jun 18 15:16:08 2009 Local Options hash (VER=V4): '48593abd'
Thu Jun 18 15:16:08 2009 Expected Remote Options hash (VER=V4):
'4b91e501'
Thu Jun 18 15:16:08 2009 Socket Buffers: R=[124928->131072] S=[124928-
>131072]
Thu Jun 18 15:16:08 2009 UDPv4 link local (bound): [undef]:1194
Thu Jun 18 15:16:08 2009 UDPv4 link remote: 173.45.227.64:1194
Thu Jun 18 15:16:14 2009 Peer Connection Initiated with
173.45.227.64:1194
Thu Jun 18 15:16:14 2009 Initialization Sequence Completed
```

Now you can *ping* the client from the server and server from the client:

```
root@vpn.unixclinic.net # ping -c2 10.8.0.2
PING 10.8.0.2 (10.8.0.2) 56(84) bytes of data:
64 bytes from 10.8.0.2: icmp_seq=1 ttl=64 time=256 ms
64 bytes from 10.8.0.2: icmp_seq=2 ttl=64 time=256 ms
```

```
root@client # ping -c2 10.8.0.1
PING 10.8.0.1 (10.8.0.1) 56(84) bytes of data:
64 bytes from 10.8.0.1: icmp_seq=1 ttl=64 time=256 ms
64 bytes from 10.8.0.1: icmp_seq=2 ttl=64 time=255 ms
```

This completes our static key-based set-up. Let us now move on to certificate-based authentication.

Certificate-based authentication

The static key-based set-up we saw earlier limits us to one client connecting to one server and also does not provide enough security. OpenVPN's support for PKI (Public Key Infrastructure) allows multiple clients to connect to a server, securely. The number of clients is only limited by the available bandwidth and the hardware resources available with the server.

If you are setting up OpenVPN for your corporate use, then you may want to use a commercial CA for issuance of certificates. However, if you want to avoid that additional cost, you can create your own CA for your internal use.

For those who want to set up a quick PKI for use by OpenVPN only, or for laboratory testing of the concept before rolling out the production set-up, OpenVPN bundles a set of scripts and tools called *easy-rsa*. You will get *easy-rsa* from any OpenVPN version that is downloaded from the main website.

If you have installed the packaged openvpn from your distribution, you should typically find this at */usr/share/doc/openvpn/examples/easy-rsa* in Debian and */usr/share/openvpn/easy-rsa* in CentOS. Please feel free to copy this directory to a convenient location of your

choice. I usually put this in the */opt/openvpn-easy-rsa-2.0/* directory and place that in my path, so I will assume that is the case here:

```
root@vpn.unixclinic.net # mkdir -p /opt/easy-rsa/keys
```

If you have a CentOS server then copy as follows:

```
root@vpn.unixclinic.net # cp -r /usr/share/openvpn/easy-rsa/2.0/* /opt/
easy-rsa/
```

On a Debian server, copy the *easy-rsa* as follows:

```
root@vpn.unixclinic.net # cp -r /usr/share/doc/openvpn/examples/easy-
rsa/2.0/* /opt/easy-rsa/
```

Next, put the directory in your path to work more easily with *easy-rsa*.

```
root@vpn.unixclinic.net # export PATH=$PATH:/opt/easy-rsa
```

The file */opt/easy-rsa/vars* contains the required environment variables for generating the keys. This file is commented very well and hence I would not describe it here. Most of the variables can be left to their default values. Make sure you change the *EASY_RSA* variable to the directory in which you've copied the *easy-rsa* tools. The default value may not suit our requirements:

```
export EASY_RSA="/opt/easy-rsa"
```

You may want to change the following variables. Given below are the values for these variables that are used for the purpose of this article:

```
export KEY_COUNTRY="IN"
export KEY_PROVINCE="UA"
export KEY_CITY="Nainital"
export KEY_ORG="Unixclinic"
export KEY_EMAIL="casupport@unixclinic.net"
```



Note: By default, the expiry of the root CA key and certificates is set to 3650 days (10 years). Practically, the certificates should expire after one or two years. So I went ahead and changed the *KEY_EXPIRE* to 365 days.

Now let us start generating the necessary keys and certificates. To start with, we need to source the variables specified in the *vars* file and make sure that the existing keys directory is cleaned. Pay extra attention to the *'* in the following command. *'* will make sure that the environment variables will be set up in the current shell.

```
root@vpn.unixclinic.net # . vars
```



Note: If you run `./clean-all`, it will be equal to an `rm -rf /opt/easy-rsa/keys/`

```
root@vpn.unixclinic.net # ./clean-all
```

Next we need to build the Diffie-Hellman (DH) parameters. It took me less than a minute to generate a 1024-bit long safe prime on my server—I have snipped the output:

```
root@vpn.unixclinic.net # build-dh
Generating DH parameters, 1024 bit long safe prime, generator 2
This is going to take a long time
.....+
.... [snipped..]
```

Next we will generate the root CA certificate and key:

```
root@vpn.unixclinic.net # build-ca
Generating a 1024 bit RSA private key
.+++++++
..+++++++
writing new private key to 'ca.key'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a
DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [IN]:
State or Province Name (full name) [UA]:
Locality Name (eg, city) [Nainital]:
Organization Name (eg, company) [Unixclinic]:
Organizational Unit Name (eg, section) []:
Common Name (eg, your name or your server's hostname) [Unixclinic CA]:
unixclinic.net
Name []:Unixclinic CA
Email Address [casupport@unixclinic.net]:
```

Now that our CA has been set up, we need to create a certificate for our VPN server. This can be done as follows:

```
root@vpn.unixclinic.net # build-key-server vpn.unixclinic.net
Generating a 1024 bit RSA private key
....+++++++
..+++++++
writing new private key to 'vpn.unixclinic.net.key'
```

You are about to be asked to enter information that will be incorporated into your certificate request.

What you are about to enter is what is called a Distinguished Name or a DN.

There are quite a few fields but you can leave some blank

For some fields there will be a default value,

If you enter '.', the field will be left blank.

```
-----
Country Name (2 letter code) [IN]:
State or Province Name (full name) [UA]:
Locality Name (eg, city) [Nainital]:
Organization Name (eg, company) [Unixclinic]:
Organizational Unit Name (eg, section) []:
Common Name (eg, your name or your server's hostname) [vpn.unixclinic.net]:
Name []:OpenVPN Server
Email Address [casupport@unixclinic.net]:vpnsupport@unixclinic.net
```

Please enter the following 'extra' attributes to be sent with your certificate request

A challenge password []:password

An optional company name []:Unixclinic

Using configuration from /opt/easy-rsa/openssl.cnf

Check that the request matches the signature

Signature ok

The Subject's Distinguished Name is as follows

```
countryName      :PRINTABLE:'IN'
stateOrProvinceName :PRINTABLE:'UA'
localityName      :PRINTABLE:'Nainital'
organizationName  :PRINTABLE:'Unixclinic'
commonName        :PRINTABLE:'vpn.unixclinic.net'
name              :PRINTABLE:'OpenVPN Server'
```

```
emailAddress      :IA5STRING:'vpnsupport@unixclinic.net'
```

Certificate is to be certified until Jun 22 04:24:13 2010 GMT (365 days)

Sign the certificate? [y/n]:y

1 out of 1 certificate requests certified, commit? [y/n]y

Write out database with 1 new entries

Data Base Updated

Now we will have to copy the server certificate to the `/etc/openvpn/keys` directory and make sure that the private key has the most restrictive permissions:

```
root@vpn.unixclinic.net # cp -v /opt/easy-rsa/keys/vpn.unixclinic.net.{crt,key} /opt/easy-rsa/keys/dh1024.pem /etc/openvpn/keys/
'/opt/easy-rsa/keys/vpn.unixclinic.net.crt' -> '/etc/openvpn/keys/vpn.unixclinic.net.crt'
'/opt/easy-rsa/keys/vpn.unixclinic.net.key' -> '/etc/openvpn/keys/vpn.unixclinic.net.key'
'/opt/easy-rsa/keys/dh1024.pem' -> '/etc/openvpn/keys/dh1024.pem'
```

```
root@vpn.unixclinic.net # cp -v /opt/easy-rsa/keys/ca.crt /etc/openvpn/keys/
```

```
'/opt/easy-rsa/keys/ca.crt' -> '/etc/openvpn/keys/ca.crt'
```

```
root@vpn.unixclinic.net # chmod 0600 /etc/openvpn/keys/vpn.unixclinic.net.key
```

Let us configure the server configuration file, which I will call *keyauth-server.conf*. Copy the contents till the line that states "Change the following for different configurations". Add the following lines to the configuration:

```
##### PKI authentication

dh /etc/openvpn/keys/dh1024.pem

ca /etc/openvpn/keys/ca.crt
cert /etc/openvpn/keys/vpn.unixclinic.net.crt

# The following file should be kept very secret.
key /etc/openvpn/keys/vpn.unixclinic.net.key

# Specifies the range of IP addresses allocated by server to client.
# The server itself will take 10.8.0.1 as its IP address.
server 10.8.0.0 255.255.255.0

# Makes sure that if available the client always gets the previous IP address.
# The record of IP addresses allocated to client is in ip.txt file.
ifconfig-pool-persist ip.txt

# Maximum number of clients which can connect, default is 100.
max-clients 10
```

Now we need to generate the client certificate as follows:

```
root@vpn.unixclinic.net # ./build-key ajitabh
Generating a 1024 bit RSA private key
...++++++
.....++++++
writing new private key to 'ajitabh.key'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a
DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [IN]:
State or Province Name (full name) [UA]:
Locality Name (eg, city) [Nainital]:
Organization Name (eg, company) [Unixclinic]:
Organizational Unit Name (eg, section) []:
Common Name (eg, your name or your server's hostname) [ajitabh]:
```

```
Name []:Ajitabh Pandey
```

```
Email Address [casupport@unixclinic.net]:ajitabhp@unixclinic.net
```

```
Please enter the following 'extra' attributes
to be sent with your certificate request
```

```
A challenge password []:password
```

```
An optional company name []:
```

```
Using configuration from /opt/easy-rsa/openssl.cnf
```

```
Check that the request matches the signature
```

```
Signature ok
```

```
The Subject's Distinguished Name is as follows
```

```
countryName      :PRINTABLE:'IN'
```

```
stateOrProvinceName :PRINTABLE:'UA'
```

```
localityName      :PRINTABLE:'Nainital'
```

```
organizationName  :PRINTABLE:'Unixclinic'
```

```
commonName        :PRINTABLE:'ajitabh'
```

```
name              :PRINTABLE:'Ajitabh Pandey'
```

```
emailAddress       :IA5STRING:'ajitabhp@unixclinic.net'
```

```
Certificate is to be certified until Jun 23 04:38:27 2010 GMT (365 days)
```

```
Sign the certificate? [y/n]:y
```

```
1 out of 1 certificate requests certified, commit? [y/n]
```

```
Write out database with 1 new entries
```

```
Data Base Updated
```

While creating the client certificate, pay special attention to the CN (Common Name); it should be unique as well as easily identifiable -- just in case you decide to do the client-specific configuration later on. Also try not to put a blank space in the CN. After the client key is created, transfer the *.crt*, *.key* and *ca.crt* files securely to the client using a pre-existing secure channel. I generally consider SCP/SFTP to be pretty safe. You can use USB thumb drives to transfer the client certificate.

Please note that the OpenVPN server does not need to know anything about the client certificates. So feel free to generate the client certificates on a separate machine, but make sure that they are signed by the same CA using which the server's keys are signed. I used the same machine to generate the client certificate.

If you are using the OpenVPN client then add the following in the configuration for PKI authentication:

```
## PKI Authentication
```

```
# Tell the OpenVPN that we are client.
```

```
# This will be use full to pull configuration settings from server later on.
```

```
client
```

```
ca /etc/openvpn/keys/ca.crt
```

```
cert /etc/openvpn/keys/ajitabh.crt
```

```
key /etc/openvpn/keys/ajitabh.key
```

Start the OpenVPN server and client, and they should connect. Check out which IP address has been allocated

by the server to the client before pinging. The server does not allocate the IP addresses in a sequence.

Using IPTABLES for simple internal network access

In a multi-client set-up, you may often need to provide protected network access to the VPN clients connecting to the server. There are two ways of doing this:

- Configure the default router of your organisation to send all traffic destined for VPN addresses (10.8.0.0/24 in our case) to the VPN server.
- Masquerade all the VPN traffic to originate from the VPN server's internal IP address (not the 10.8.0.1).

In most cases, my clients have chosen to go for masquerading as they host their server in an external data centre and do not have control over the router. This can be done as follows:


```
root@vpn.unixclinic.net # iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
```

I am assuming that the VPN server is connected to the internal network through the *eth0* interface.

Further, one of my clients had the requirement of forcing all port 80 traffic coming from the VPN clients to go through their Squid proxy. This can be achieved by following the simple IPTABLES rule:

```
root@vpn.unixclinic.net # iptables -t nat -A PREROUTING -i tun+ -p tcp --dport 80 -j REDIRECT --to-port 3128
```

Moving on

In this article we have looked at two very commonly used configurations for OpenVPN. I have used OpenVPN with various types of configurations such as authentication using a MySQL database, freeradius backends, PAM, etc. If time permits, I will try to cover these configurations in some future article. 

References

- OpenVPN documentation: <http://openvpn.org/index.php/open-source/documentation.html>
- OpenVPN FAQ: <http://openvpn.org/index.php/open-source/faq.html>
- OpenVPN articles: <http://openvpn.org/index.php/open-source/articles.html>
- The user-space VPN and OpenVPN: <http://openvpn.net/papers/BLUG-talk/index.html>

By: Ajitabh Pandey

The author has more than 13 years of diversified IT industry experience in training, support and consulting. You can learn more about him at ajitabhpandey.info and e-mail him at [ajitabhpandey \[at\] ajitabhpandey \[dot\] info](mailto:ajitabhpandey[at]ajitabhpandey[dot]info)

LINUX

ForYou

Calling...

Share your
knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of linuxforu.com we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developers or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools





One Extra Step, Courtesy GCC

Compilers fill the gap left by the limited set of constructs in C. We learn some GCC-related tricks in this article.

Being a developer, I'm fascinated by well-structured programs. Although the C programming language has a limited set of constructs, and Dennis Ritchie apparently has attempted to touch everything, sometimes we stumble upon an amazing piece of code, which makes us go: "I wonder if I can write it in C, too!!!"

However much smart code you write, with C being blessed with a limited set of constructs, compilers always tend to cross the boundaries there and provide a lot more to the programmers—giving additional functionality, short-hand programming and better performance.

Here we learn some of what is provided by GCC.

Identify the types

The RTTI (Run Time Type Identification) is a feature of the C++ language wherein you can play with the 'types' of variables to create interesting things. This could not be achieved staying within the limits of the C language. However, something of this sort can be attempted at compile time. Consider a simple macro to swap two integers:

```
#define SWAP(a,b) {int temp; temp = a; a=b; b=temp}
```

The same macro could be made available for floats, chars, and even structures and unions by using the following command:

```
#define SWAP(a,b) { typedef (a) temp = (a); a=b; b=temp}
```

Inline functions

Implementing a functionality as a macro or a function is always a trade off. Hence, it is always preferable to leave the decision to the compiler itself. The keyword *inline* before the function declaration leaves the decision to GCC (depending upon whether you choose to optimise for size or for performance) when you compile the code.

Thread local storage

GCC, in conjunction with modern thread libraries, supports a special extension wherein variables could be declared local to a thread. The way to declare such variables is to use the keyword `__thread`. For example, the following statement declares an integer in a given thread context:

```
__thread unsigned int Sample;
```

This feature becomes extremely useful when writing multi-threaded applications.

Note that since 'Sample' is a thread local variable, the address of the variable is evaluated at run-time and the pointer arithmetic on it would lead to unexpected results.

Designated initialisers within arrays

Consider the following declaration:

```
enum
{
    STATE_COMPRESSION,
    STATE_EXPANSION,
    STATE_EXHAUST
};

char * states[] =
{
    "Compression",
    "Expansion",
    "Exhaust",
};
```

Hence,

```
printf("%s", elements[STATE_COMPRESSION]);
```

...will output to:

```
Compression
```

But, tomorrow if we end up swapping the position of the *STATE_COMPRESSION* and *STATE_EXHAUST* in the enumeration, the names of the objects will change and the output of the same *printf* statement will be:

```
Exhaust
```

Here is how we can write the same with GCC:

```
char * states[] =
{
    [STATE_COMPRESSION] = "Compression",
    [STATE_EXPANSION] = "Expansion",
    [STATE_EXHAUST] = "Exhaust",
};
```

And come what may, the *STATE_COMPRESSION* is *Compression* and *STATE_EXHAUST* is *Exhaust*.

Case range

Short programs are always more readable and smarter. Consider the following code snippet:

```
switch(value)
{
    case 1:
    case 2:
    case 3:
    case 4:
        printf("First four");
        break;
    case 5:
    case 6:
    case 7:
        printf("Till seven")
        break;
    case 8:
    case 9:
        printf("Big bros")
        break;
    default:
        printf("Too much.....!!!");
```

Well, there's no need to write so much. You can simply get away with:

```
switch(value)
{
    case 1 ... 4:
        printf("First four");
        break;
    case 5 ... 7:
        printf("Till seven")
        break;
    case 8 ... 9:
        printf("Big bros")
        break;
    default:
        printf("Thanks GCC.....");
```

Variable length arrays

A typical problem in programming is the lack of knowledge of limits—particularly the paths and filenames. GCC comes up with a simple and elegant solution for it. Consider the following code snippet:

```
FILE* OpenFile( const char *path, const char* filename, const char*
mode)
{
    char buffer[strlen(path) + strlen(filename) + 1];
    strcpy(buffer,path);
    strcat(buffer,filename);
    return fopen(buffer,mode);
}
```

The above works seamlessly on GCC and produces the expected results.



Note: A more portable way of doing this is to use the *alloca()* function. The syntax is similar to *malloc*, but allocates the memory on a stack frame of the caller. The memory is released when the caller function returns.

Optimisation on likelihood

Consider the following function:

```
extern unsigned int g_event; /* Set to true if event has occurred*/
unsigned int my_poll()
{
    static int OccurCount = 0;
    static int FailCount = 0;
    if(g_event)
    {
        /* Handle the event*/
        OccurCount += 1;
    }
    else
    {
        /* Handle if the event has not occurred */
        FailCount += 1;
    }
}
```

Now the code for *if(g_event)* in the assembly with *-O2* optimisation looks like the following:

```
a: 84 c0      test  %al,%al
c: 74 12      je    20 <_my_poll+0x20>
```

So the compiler assumes that the value of *g_event* is likely to be non-zero and generates a branch instruction when *g_event* is zero.

This adds a branch penalty if the *g_event* is unlikely and the *else* part is the usual path. This could be communicated to the compiler as follows:


```
#define likely(x)    __builtin_expect((x),1)
#define unlikely(x)  __builtin_expect((x),0)

extern unsigned int g_event; /* Set to true if event has occurred*/
unsigned int my_poll()
{
    static int OccurCount = 0;
    static int FailCount = 0;
    if(unlikely(g_event))
    {
        /* Handle the event*/
    }
```

And we tell the compiler that *g_event* is a rare event. So the assembly changes to:

```
a: 84 c0      test  %al,%al
c: 75 0b      jne   19 <_my_poll+0x19>
```


This achieves a great improvement in performance.

It will be good to note once again that all these little functionalities offered by GCC are not standard. The code written in such a way is not guaranteed to work on another C compiler. But as long as we stick to GCC, all this will definitely work. Most of these functionalities are even used extensively in the Linux kernel code. So why can't we use them, too? 

By: Nilesh Govande

The author is a Linux enthusiast and can be contacted at nileshgovande@yahoo.com. His areas of interest include Linux systems software, application development and virtualisation. He is currently working with the LSI Research & Development Centre, Pune.


THE COMPLETE MAGAZINE ON OPEN SOURCE



Your favourite Linux Magazine is
now on the Web, too.

LinuxForU.com

Follow us on Twitter @LinuxForYou



Defect Prevention and Identification

In this article, we will look at the various practices and techniques that can help you prevent defects in your software, and how to catch them if they already exist.

*M*any of us know that following a proper procedure during the software development cycle should give us close to zero-defect programs that meet users' requirements. The question is, how to go about this?

When customers (the users who require the software we develop) provide you with their requirements, it's not that you start working on them there and then. You first need to understand clearly what the customers want. Once you have gone through the requirements, put down what you have understood. Then, get your understanding confirmed from the customer. Doubts, if any, in the requirement specifications, must be clarified at this stage.

Do not procrastinate or hesitate in asking your questions. The diagram in Figure 1 pictorially represents the procedure.

While working on the development of systems software tools, I have seen many defects creeping in due to the misinterpretation of the requirement specifications. Fixing such defects at a later stage can prove costly. So, it is very important that you get your understanding verified from the customer before you start working on the code.

Peer review

Functioning as a team is a skill. Delivering a high-quality product is not the responsibility of a single individual. It's the entire team



Illustration: Pradeep Kumar

who is responsible for it. If the product fails, each team member is responsible for it.

Peer review is an important part of team work. In this process, a team member requests an artefact to be reviewed. The other team members then provide their review comments, which may include corrections, suggestions and doubts about the artefact. The artefact is then updated based on the review comments. This process is repeated till the artefact is up to the satisfaction of all the team members. Of course, in case of any conflict, it is the project leader who makes the final decision. Figure 2 gives a pictorial representation of the procedure.

Reviews are critical at the following stages of software development:

1. Requirement-specification review: As discussed earlier, understanding the customers' requirements is very important. So, while you prepare the requirement specification, get it reviewed by your team members. Ten people in a team may have




Figure 1: Analysing application users' requirements

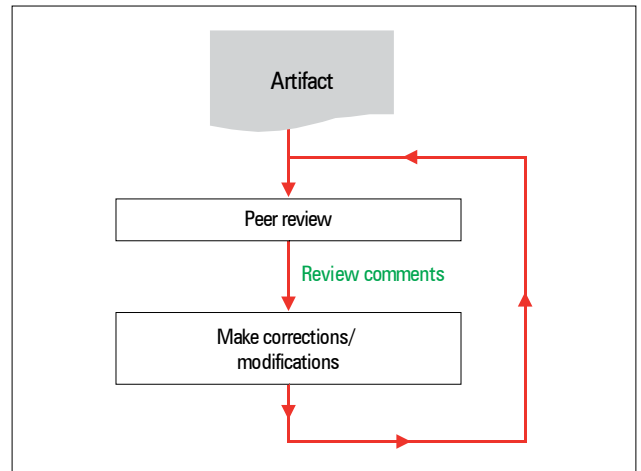


Figure 2: Steps involved in peer review

ten different interpretations of the specifications. Discuss it within your team, and then pass on your understanding to the customer for verification.

- 2. Design review:** Once the requirement specification gets finalised, we move on to the design phase. In the design stage, you would now think about how to approach the problem. As you will agree, review at this stage too is vital. The selection of the wrong strategy can put the entire system in a miserable state. Reviews performed at this stage will help you in:
 - Analysing various strategies to solve the problem;
 - Analysing the feasibility of each strategy;
 - Knowing the advantages/disadvantages of each strategy.
- 3. Code review:** This involves examining the source code to spot defects. While working on the development of systems software tools, I have been a witness to how code review can really help you find defects. It is good to let someone in your team walk through your code, and do the review. All the members must review the code changes with respect to other modules and give their feedback in case some side-effect is suspected. Many defects like memory leaks, wrong passing of arguments, unreachable code, lack of readability, high complexity and maintainability issues can be identified via the code review. Finding defects at the coding stage, and fixing them immediately, would prove to be less expensive than finding them in the testing stage.

Source code documentation

Maintaining software is a crucial aspect of the software development process. Poorly documented, or rather undocumented code, is quite difficult to maintain. Having somebody modifying your undocumented code can result in disastrous consequences. So, it's important to document your code and the assumptions made by you, while coding.

Keep the documentation close to the source. Multiple

documents are difficult to maintain. A tool such as *doxygen* is useful for generating source code documentation.

Use a coding checklist

A self-review coding checklist must be maintained by the developer. The checklist should be updated as and when defects are found. And it should be updated with actionable checklist points, which could ensure that any defect is not introduced. Using the checklist would ascertain that the same defect is not repeated. All members in the team should use that checklist and update it periodically.

Analysing the ‘5 Whys’

Despite the quality processes you adopt, you must admit the fact that defects do somehow creep in. After all, software developers are only human. So when you find a defect, what should you do with it? Most software developers would be tempted to ‘fix and close it’.

But isn't it important to find out what caused the defect? Yes, it is. Analysing the cause of a defect can let you catch some other defects that may exist, and prevent similar defects being introduced in the future.

I have found the ‘5 Whys’ method [as described in the book *Kaizen: The Key to Japan's Competitive Success* by Masaaki Imai (www.amazon.com/Kaizen-Key-Japans-Competitive-Success/dp/007554332X)], to be extremely useful to reach the root cause of the defect.

The ‘5 Whys’ method involves the asking of questions to identify the root cause of a problem. For example, if a person dies due to an accident, the objective is to identify why he died. So, we put up a chain of questions, until we reach the root cause of the problem. This is done as follows:

John died. [The Problem]

1. Why? – He met with an accident.
2. Why? – He was driving his car quite fast.
3. Why? – The car brakes were not working.
4. Why? – The brake oil was not checked for a long time, and hence brakes failed. [The root cause has been found in the ‘4th Why’, itself]

There need not always be five questions to be asked. Sometimes, you may need to keep on questioning until you reach the root cause. And in some cases, you might end up reaching the root cause in just three questions! The idea behind this method is to drill-down to reach the root cause of the problem.

Knowing the root cause of the defect can help you to prevent similar defects in future.

Frequent releases

You need to plan for frequent releases, because releases at regular intervals maintain the tempo of the team. At the same time, you get a chance to refine your processes based on the feedback and the lessons learnt. It would also ensure that there is no slippage in the schedule.

Tests and testability

Write code that is testable. If your code is not testable, it implies complex code, and hence poor maintainability and higher chances of defects.

Do the testing, starting with the assumption that there are defects. Take different input combinations. Take not only the positive patterns, but the negative patterns as well. List down the following when you create a test case:

- Purpose of the test case [what it intends to test]
- Expected output
- Hardware/software requirements

For developers, it is highly important to test whatever they code. Performing unit testing is an effective method to find defects in your code. See to it that your code can be isolated from the rest of the code for unit testing.


In a team, we have several developers working on different modules. The challenge is to seamlessly integrate all the changes to form a single product. Often, we fail to analyse the impact of changes in one module with respect to other modules. What I have experienced is that many defects creep in due to interface problems. Defects come in when an interface is either not tested at all, or is not tested fully. So, the interface must be tested while integrating changes in a module with the rest of the modules.

While you test, check whether your code is doing what it is supposed to do. More importantly, check whether it is doing what it is not supposed to. If it is doing something it should not be doing, that's a defect.

Some handy tips

By catching defects early, the amount of time spent in fixing things later is reduced. Keep the following steps in mind to prevent defects, and catch them as early as possible:

- Avoid introducing new defects in the code.
- Use a checklist.
- Code should be simple, as simple code is easy to maintain.
- Avoid global variables and magic numbers.
- Complex design leads to poor testability.
- The user should be central to the design.
- Test as you code.
- Change of code should support ease of testing, debugging and modification.

As in everything else, prevention is better than cure. So prevent the defects in your program, and you won't need to waste resources on finding them. 

By: Deepti Sharma

The author is a senior software engineer (team lead) at Acme Technologies. She has considerable amount of experience in the development of compilers and tools for embedded systems. She can be reached at [deepti.DOT.acmet AT gmail DOT com](mailto:deepti.DOT.acmet@gmail.com).

Python in Research Computation with Polynomials



Let's explore some of the common computational needs and how they can be met using Numpy.

Numpy is useful when complex mathematical computations are involved. This includes engineering problems for which Matlab, a very expensive product, is often used. See www.scipy.org/NumPy_for_Matlab_Users for more details. In this article, we will explore some common computational needs and look at how Numpy can address them.

Regression analysis / least square fit

A very common requirement is to have a model of a lot of collected data. You need to test whether the data fits the model. In this article, you can explore the Galilean model because, as mentioned in

www.vpri.org/pdf/rn2005001_learning.pdf, "This subject has been extensively studied by college students in the US: 70 per cent (including science majors) fail to understand this Galilean model of gravity near the surface of the Earth." It roughly means that a large proportion of the population believes that heavier objects fall faster. With the technology available these days, experimental testing of these beliefs is not difficult.

(Just for fun, check out the hammer and feather drop—Apollo 15 astronaut, David Scott, demonstrates on the Moon that Galileo was right, at en.wikipedia.org/wiki/Galileo%27s_Leaning_Tower_of_Pisa_experiment)

You can take a video of a few falling objects and convert them into frames. So you may get about 30 frames per second. Identifying an object in a video frame will not be very precise but the errors in measurement are an inherent part of research. The model you wish to verify is: $y = \frac{1}{2} g t^2 + v_0 t + y_0$

You have a list of the object's position every 1/30th of a second. You can fit the values to a second-degree polynomial least-square fit algorithm. Here is the code you would use:

```
import numpy as np
values = [195.7,193.51,190.59,185.84,180.73,173.43,
165.76,157.36,147.87,136.92,124.87]
y = np.array(values)
t = np.linspace(0, 10.0/30, 11)
poly_coeff = np.polyfit(t, y, 2)
print poly_coeff
print 'g = ', 2*poly_coeff[0]
```

You convert a Python list of 11 values into a Numpy array. The *linspace* method creates an equally spaced array, with the first value being 0 and the last (11th) value being 1/3 sec. The *polyfit* method will fit the data to a polynomial with the degree being 2 and return an array of three numbers. These three numbers are the coefficients of the polynomial starting with the highest power first. So, the first value is half of *g*. The above data gives *g* as -963 cm/sec². The negative sign indicates the downward direction.

So, how good was the fit?

Before discussing that, let us look at the relationship between an array and a polynomial.

Polynomials and arrays

Any polynomial can be represented by a vector, just by keeping track of the coefficients. For example, [1,0,-4] would represent ($x^2 - 4$). Try the following examples:

```
>>> import numpy as np
>>> p2 = np.poly1d([1,0,-4])
>>> p2(0)
>>> p2([0,1,2])
```

You have created an object of the type polynomial (or simply a polynomial function) of degree 2. Now, you can use it like you would a function. The result for *p2(0)* would be -4 and for *p2([0,1,2])*, it would be the array [-4, -3, 0].

The best way to see how well the polynomial you obtained fits the data is to visually see it through a plot.

Matplotlib

The *python-matplotlib* package must be installed.

You will need to plot the original data and the values computed from the polynomial. Hence, add the following




Figure 1: A Falling Object - Observed positions and the theoretical fit to a quadratic polynomial

code after obtaining the least square fit to your data:

```
import matplotlib.pyplot as plt
plt.plot(t, y, 'o')
plt.plot(t, np.poly1d(poly_coeff)(t), '-')
plt.show()
```

The first call to the plot function plots small circles at the actual pairs of data points (t_i, y_i).

In the second call to the plot function, you are creating a polynomial function from the array *poly_coeff* and then computing the values of this function for each value of *t*. Finally, you are plotting these values against *t* as a line. The resulting fit for the above example is pretty good as can be seen in Figure 1.

The roots

Solving a polynomial is another common requirement. The roots of the polynomial that fit to the data would indicate when the object would be on the ground. Add the following lines to your code:

```
roots = np.roots(poly_coeff)
print 'Roots', roots
```

The output of the program, including the earlier code, would be:

```
[-481.72027972 -52.31748252 195.86951049]
g = -963.440559441
Roots [-0.69426607 0.58566055]
```

The meaningful root in this case is .585 sec. It may not be very useful in this context, but it does become valuable when dropping food packets or bombs from an aeroplane.




Figure 2: A plot showing observations and polynomial fits for two falling bodies

A polynomial can be represented by the coefficients as we have discussed earlier. It can be equally well represented by its roots. So, if you create an object of type polynomial, both representations are easily available to you, as follows:

```
>>> pf=np.poly1d(poly_coeff)
>>> print 'The roots', pf.r
The roots [-0.69426607  0.58566055]
>>> print 'The coefficients', pf.c
The coefficients [-481.72027972 -52.31748252 195.86951049]
```

Multiple data sets

Typically, there will be multiple runs and you will need to analyse each set of data. In the following example, you have two runs of data—one for a stone and the second for a small plastic cap. Try the following code:

```
import numpy as np
set1 = [195.7,193.51,190.59,185.84,180.73,173.43,
        165.76,157.36,147.87,136.92,124.87]
set2 = [194.04,192.6,189.36,185.04,180.36,173.52,
        167.76,158.76,148.68,138.6,127.8]
y = np.array([set1,set2])
print 'Shape of the array', y.shape
t = np.linspace(0, 10.0/30, 11)
res = np.polyfit(t, y.transpose(), 2)
print 'The shape of the result', res.shape
print 'g = ', 2*res[0]
```

The array *y* is the list of data values for each set. However, *polyfit* requires the list to be grouped by corresponding values for all trial data sets. So, you need to take the transpose of the array *y* to fit the needs of the *polyfit* function. The result will be a set of three coefficients, one for each set of data. You get a pair of

values for *g*. The output will now be as follows:

Shape of the array (2, 11)

The shape of the result (3, 2)

```
g = [-963.44055944 -953.11888112]
```

You may need to do more experiments with plastic caps! But meanwhile, you can plot the two sets of data and the fitted lines as shown below:

```
import matplotlib.pyplot as plt
plt.plot(t,set1,'o', label='Dataset 1')
plt.plot(t,set2,'d', label='Dataset 2')
coeff = res.transpose()
plt.plot(t,np.poly1d(coeff[0])(t),'-', label='Fit 1')
plt.plot(t,np.poly1d(coeff[1])(t),'-', label='Fit 2')
plt.legend()
plt.show()
```


Pyplot will choose a different colour for each plot. The first set of points is drawn as small circles. The second set of points is drawn as little diamonds. The last two are continuous lines. Note that you needed to take the transpose of the result matrix for easy access to the coefficients of each polynomial. A legend will be useful in this case. See Figure 2.

You can find the details of all the options of plotting by using inbuilt help for Python:

```
>>> help(plt.plot)
```


What else

Polynomials are very useful. The integral of a polynomial is another polynomial. The derivative is another polynomial. Numpy includes functions, *polyint* and *polyder* for this purpose. The function *polymul* will return another polynomial that is the product of two polynomials. There is a wealth of functions available in Numpy, which makes working with polynomials in education and research a much easier task.

The Scipy module of Python includes everything of Numpy and some more. So, you can replace the import statement by '*import scipy as np*' and your code will still run fine. In case you are still hesitating to use Python for science, the article 'Python in Science: How long until a Nobel Prize? And Now For Something Completely Different' by Doug Hellmann, might interest you—www.doughellmann.com/articles/CompletelyDifferent-2007-11-science. **END** 

By: Dr. Anil Seth

The author is a consultant by profession and can be reached at seth.anil@gmail.com



The Programmable Web

Have you ever wondered what else you could do on the Web? Whether a simple mail of yours could execute a chain of programs on the Web? The answer to your doubts lies in the programmable Web, also known as WebHooks.

I have been programming in various script languages for four years now and there have been instances when I created many cascaded requests. Cascaded requests are different in a way that a single request is responsible for performing more than one task. Some gurus also call them chained requests.

The programmable Web is a decentralised concept that is not owned or governed by a company. This means that even you can contribute to it. People across the globe are trying to understand it more deeply and some Internet majors have already implemented it. In simple words, it is a Web where you, as the user, can 'pipe' data between applications, much like the UNIX command line.

Google, being one of the leaders in adapting the programmable Web, introduced WebHooks in many of its applications like Google Alerts, Google Reader and Feedburner.

What are WebHooks?

WebHooks let you extend Web applications, customise them and integrate them using something that cannot be accessed by programming. For a Web developer, it is another design pattern that has the ability to make Web requests and store some extra data about the users. From a user's point of view, they are applications that can generate/trigger events to get data in real-time from that very Web application.

Why do we require WebHooks?

Chaining: They allow your applications to work as a chain or hook linking each to the other, more like executing a program that has a loop, some conditions, etc.

Notifications: When you receive an e-mail, the notification reads, "1 new e-mail received." What if you want the notification to do something more than that?

Data synching: How many times have you had to import your contacts into Gmail, then to Thunderbird or to some other Web

application? Using WebHooks, you can share data among different applications to reduce, if not eliminate, redundancy.

Extensibility: WebHooks let your users extend your applications to build plug-ins on their own. All you have to do is just specify the call back URL.

And all this can be achieved with the new design that allows users to specify their callback URLs in the application. The application can then return the data to that URL. This way, the calling user does not require access to the application's source code. It's more like an API but in an API call, the user has to invoke the request explicitly, and if there is a WebHook associated with the application, then the callback URL gets notified automatically without having to request the API.

You might find some examples in your daily routine when you use certain applications that seem to work like a WebHook. But they are not the real-time examples. Also, the scenario mentioned above is just a simple example. The practical requirement may vary and this is what triggers the necessity to start using WebHooks.

How do WebHooks work?

A consumer of a Web service or a user is provided with URLs for different events, to which an application will post data when the events occur. Here's an example: on receiving an e-mail, the mail program will be clever enough to post the new message to some URL (<http://mydomain.in/webhooks.php>). So you can avoid repeatedly checking your inbox or pressing the send/receive button of your mail client. The event gets fired as soon as the mail is received. Now it's up to you how to utilise that event. You can keep SMS alerts on that event, or you can write a script to perform some other operation when you receive that e-mail.

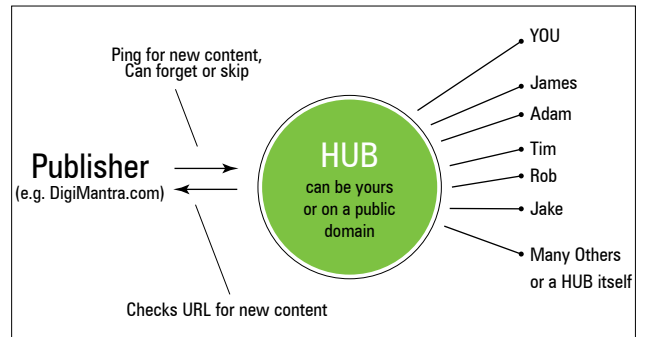
An example to get you started

PubSubHubBub

Have you heard the term PubSubHubBub? I am sure many of you have not. In simple words, it's a new kind of RSS and a step closer to the programmable Web. Some of you might wonder why we need a new RSS? Allow me to dig a little deeper.

How RSS works

RSS, as we all know, is content syndication and is served in the form of a XML file to a RSS reader. Programmers and many Web applications use these feeds to drive many powerful applications like blog aggregators and mashing services. But the problem with RSS is the repeated polling. Polling is a process in which a client pings the server to check for new content. RSS clients check servers for new content after a regular interval of time, set by the user. The server has to respond to every request saying, "You have got the latest content." You can change the ping intervals to the server, but there is no ideal time for it. If you set it to one minute, then the server is pinged every minute, causing overheads to



both the server and the requester. And if the delay is of one hour or more, then it's too late. And it is impossible to set the time correctly, even with an adaptive rate control. However PubSubHubBub tries to eliminate this problem and updates you with the latest content within seconds of the server getting updated. In a scenario like this, even a minute's delay seems to be a lot.

How does PubSubHubBub work?

PubSubHubBub tries to achieve a programmable Web. Instead of repeated polling, in the normal RSS, a URL declares its hub server(s) using the `<link rel="hub"/>` tag. The hubs can be on any public domain or even on your personal server.

After subscribing, if the atom file of a RSS link declares its hub(s), the user can then avoid repeated polling of the URL and can instead register with the feed's hub(s) and subscribe to updates.

Now, the moment new content is published, the publisher pings and notifies the hub about the update. The hub fetches the updated content and broadcasts it to all the clients/users who are subscribed to it. It is more like a push protocol similar to the push e-mail service from RIM's Blackberry. Thus, the subscriber to that hub gets the new content automatically without having to ask for it, saving a lot of resources.

What is wrong with a delay of a minute or an hour?

Since feed updates are important sources of information, in the world of the real-time Web, even a few seconds sounds too long. Real-time applications are something that are achieved as part of a chain of applications. Google Wave is a very good example of the programmable Web. Many mailing services like Mailchimp have already started supporting WebHooks in their APIs. Twitter is also going to support this sometime soon.

Implementation

The implementation of PubSubHubBub requires the following entities:

- 1) Publishers
- 2) Hub(s)
- 3) Subscribers

Publishers: A publisher is the one responsible for

producing content for a blog, a regularly updated website, or any feed source.

Hub(s): A hub is the heart of this protocol. A publisher can update a single hub or multiple hubs when it publishes new content. A hub, in return, broadcasts the new content to authenticated subscribers using the push protocol methodology.

Subscribers: Subscribers are the ones who subscribe to a hub. This is generally the same as subscribing to an RSS feed. The only difference is that if the Atom file contains the definition of a hub, the RSS reader client subscribes to the specified hub instead of normal Atom feeds. This is when the latest content is served to all the users instantly.

Steps to submit

Follow the steps to announce a hub on your website, so that the users subscribing to your site's RSS feed may also be able to subscribe to the hubs.

- 1) Add a `<link />` tag in the `<head></head>` section of your HTML page, which has the address of the hub to which the latest content is going to be announced. For this example let us take the hub <http://pubsubhubbub.appspot.com> by Brett Slatkin.

```
<link rel="hub" href="http://pubsubhubbub.appspot.com" />
```

- 2) You need to mention the hub definition into the RSS feeds, so that users who have already subscribed may also get the notification of the hub.

```
<atom:link rel="hub" href="http://pubsubhubbub.appspot.com" />
```

- 3) Now that the announcement is done for the hub, we need a logic that pushes new content to the hub(s). Google code has a library that facilitates publishing your content to a single hub or multiple hubs. Download the PHP library from Google code (<http://tr.im/pubsubhubbub>). Use the following PHP code to publish it to the hub.

```
<?php
require_once(publisher.php);

//hub is defined on the following url
$hub_url = "http://pubsubhubbub.appspot.com/publish";


//create object for the class Publisher, defined in publisher.php
$publish=new Publisher($hub_url);

//URL on which new content is posted
$url="http://www.digimantra.com";

//now publish the new content using publish_update() function
$response=$publish->publish_update($url);

if($response)
    echo "Topic published, successfully";
```

```
else
    echo "Oops !! Something is broken";
?>
```

So far, we've learnt some key aspects about the programmable Web, and WebHooks is just another way of mastering them. This protocol is decentralised and is often referred to by many other terms. We should not get confused with the terminology because the motive of each term is to progress towards a real-time Web. Everyone can contribute to the development of the real-time Web by writing applications that support WebHooks. There are many examples available online, explaining the implementation of WebHooks. Google Wave is the most talked of example of the programmable Web. A Google robot built in Python is one good example to help you understand how WebHooks work, practically. To get started and build your own Google Wave robot, you can go through <http://links.digimantra.com/get-started-wave>. 

By: Sachin Khosla

The author is the founder of an emerging technology blog, Digimantra.com. He loves to code, blog, participate in community events and is a Technology Evangelist. Also known as Realin, online, he is an active promoter of modern technologies and social media.

www.eleb2b.com

Electronic Components Products Machinery

Will now be sold here

Presenting Electronics Bazaar,
India's First B2B Sourcing Guide
For Electronics Components,
Products and Machinery.

Unit Price: **Rs 50**


Subscription Rate: **Rs 350 (12 Issues)**

EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area,
Phase-I, New Delhi 110 020; Ph: 011-26810601-03; Fax: 26817563;
Email: info@efyindia.com; Website: eleb2b.com



EveryWhere & AnyWhere

32 lakh active young sports loving men read
cricket today every month*



Lifeline of Cricket

Cricket
today

www.crickettoday.com




Available in English & Hindi

*SOURCE IRS 2008

CPUID

The x86 Processor Identification Tool



With CPUID assembly instruction, software can determine information about the processor—the vendor, model, family, supported instruction sets, and other extended features. This article will try to demonstrate the use of CPUID instruction with assembly and C, on Linux.

Try peeking inside the `/proc/cpuinfo` file:

`cat /proc/cpuinfo`

...and any Linux console will show up details about the processor -- its vendor details, model, family, and capabilities. All this valuable information is available thanks to the CPUID (CPU IDentification) instructions set available in x86 processors.

This article is all about getting our hands dirty with CPUID.

CPUID: How things work

CPUID instruction is used to return the capabilities and configuration information

of x86-based processors. But how can we use CPUID to fetch this information?

Well, fundamentally the information is accessed by loading a specific value in the EAX register and executing the CPUID instruction. The CPUID instructions will then perform a specific function based on the input given to EAX. The results are finally returned to registers EAX, EBX, ECX and EDX.

- Input supported value → EAX
- Execute instruction → CPUID
- Result(s) → EAX, EBX, ECX, EDX

But why do we need CPUID? Well any operating system and application software can exploit the CPUID instruction to determine if it is compatible with the

hardware platform. It can also choose the proper execution path or run-time libraries to leverage specific processor features or do anything associated with the processor information.

Standard and extended functions

The CPUID instruction supports two sets of functions: standard and extended functions.

For standard functions -- the highest acceptable value (or range) for the EAX register input and CPUID functions that return the basic processor information—the program should input 0000_0000h to the EAX register.

```
mov eax, 00h
cpuid
```

For extended functions—the highest acceptable value (or range) for the EAX register input and CPUID functions that return the extended processor information—the program should input 8000_0000h to the EAX register.

```
mov eax, 80000000h
cpuid
```

After the execution of the above CPUID instructions, the result will be present in the EAX register.

(For more detailed information about standard and extended functions, readers can refer to the Intel Architecture Software Developer's Manual, or the AMD64 Architecture Programmer's Manual.)

CPUID with the GCC inline assembly function

C professionals/hobbyists can exploit CPUID instruction with the GCC inline assembly function, that is, `asm`.

```
__asm__ ("cpuid"
: "=a" (result[EAX]),
  "=b" (result[EBX]),
  "=c" (result[ECX]),
  "=d" (result[EDX])
: "a" (input_eax)
);
```

The GNU C compiler uses the AT&T syntax for assembly coding. To learn more about the inline assembly feature provided by GCC, you may refer to www.ibiblio.org/gferg/ldp/GCC-Inline-Assembly-HOWTO.html.

Now let's try out something more significant with CPUID using the inline assembly function.

The CPU vendor string, however, may not be extremely valuable to most developers, but it is easy to demonstrate how CPUID instruction fetches information with just a basic example.

With an input value of 0000_0000h in EAX, in addition to returning the largest standard function number in the EAX register (as discussed above), the CPU vendor ID

string can also be verified at the same time. We can see how, with this example:

```
#include <stdio.h>
#define EAX 0
#define EBX 1
#define ECX 2
#define EDX 3


int main()
{
    unsigned int result[4], input_eax;
    input_eax = 0x00;
    __asm__ ("cpuid"
: "=a" (result[EAX]),
  "=b" (result[EBX]),
  "=c" (result[ECX]),
  "=d" (result[EDX])
: "a" (input_eax)
);
    return 0;
}
```

In the above C code, after the `__asm__` function executes the CPUID instruction, `result[EBX]`, `result[ECX]` and `result[EDX]` (when combined) will represent the complete vendor ID string. (Readers should easily be able to run the above example on any GCC compiler.)

On my Intel processor-based laptop, the vendor-ID string returned is "GenuineIntel" as these registers contain the following values:

- EBX → 0x756e6547 → "ueG" ('G' in BL),
- EDX → 0x49656e69 → "Ieni" ('i' in DL),
- ECX → 0x6c65746e → "letn" ('n' in CL).

When the CPUID instruction is executed with 1 in EAX, the version and feature information is returned back to EAX.

For detailed information about processor identification and supported features (standard and extended functions), readers may refer to the Intel Architecture Software Developer's Manual or the AMD64 Architecture Programmer's Manual. 

References

- AMD CPUID Specification, Rev. 2.28: www.amd.com/us-en/assets/content_type/white_papers_and_tech_docs/25481.pdf
- Intel Processor Identification and the CPUID Instruction: www.intel.com/Assets/PDF/appnote/241618.pdf
- Intel 64 and IA-32 Architecture Software Developer's Manual Volume 2A: *Instruction Set Reference, A-M*: <http://www.intel.com/assets/pdf/manual/253666.pdf>

By: Vivek Kumar

The author is a system engineer at IBM India, and works for IBM-ISS (Internet Security System). He is based in Pune and can be reached at mail_DOT_vivekkumar_AT_rediffmail_DOT_com

A Voyage to the Kernel



Part 17

Segment: 3.6, Day 16

We have covered the details concerning interrupts and their descriptions during the previous day of our voyage.

But one point I deliberately missed was that interrupt handlers form only the first half of the interrupt processing method. The main issue with handlers is that they run asynchronously and may even interrupt with other codes (sometimes even critical ones!). The best way to avoid this is to run the handlers as quickly as possible. You can understand this point more clearly if you can visualise the way in which they deal with your hardware. For this, you may consider handlers as part of a whole mechanism set up to manage actual hardware interrupts. So we can use them for all such time-critical activities.

But we need to route the 'less critical' part to another portion where interrupts are enabled. Thus, managing the interrupts is two-fold. In the first segment (let's call that as 'top half'), the handlers are executed by the kernel asynchronously (as mentioned before) as a response to the hardware interrupt. When it comes to the second part, we deal with actions, which are linked to interrupts that are left out by the handler.

Technically speaking, this second part holds the lion's share of the whole work (since we are giving only 'quick' works to handlers). But the

handler does an important job—acknowledging the receipt of interrupts and moving data to/from the hardware. All this is covered in the 'top half'.

Let's take an example. Assume that you want to transfer some data from hardware into memory. This can be handled in the top part and the processing can be done in the bottom part.

This is an important feature as far as the programmer is concerned. This enables the device manager coder to divide the work so as to get the best results. Here are a few guidelines you can use while you write your driver:

- If the work is time-specific (say, you want to finish it soon), use a handler for the work.
- If it is hardware-specific, priority should again be given to the interrupt handler.
- If you wish to handle the processing of any data, then consider the bottom half.



Tip: Before you code your own driver, you are advised to take a look at existing interrupt handlers and bottom halves. And the key point you need to remember is: *the quicker the (handler) execution, the better.*

By using the bottom part, you can limit the work that you intend to do in the handlers since they run with the current interrupt line disabled on all processors. And in the worst case, those

that employ `SA_INTERRUPT` will have all local interrupts disabled! Hence, you can see that reducing the time allocated for this (that is, the time for which the interrupts remain disabled) is vital when it comes to the overall performance of the system.

The logical part is very vital when you decide when you want to draw the line that separates the top and bottom half (terminologically it should be *part* instead of half). The whole point of this division is to improvise the system performance. You can see that by this division we can actually postpone many works. You can perform them when the system is 'less busy'. In most cases, these bottom halves run just after the interrupt returns. This separation will enable us to have the best system performance.

There are various mechanisms that allow you to implement the bottom part effectively. Historically, Linux offered the 'bottom half' (BH) for meddling with all the bottom parts. The original interface was quite simple and somewhat elegant, providing a statically created list of 32 bottom halves. The top part set the bottom to run by assigning a bit in a 32-bit integer. Each BH was globally synchronised and no two BHs were allowed to run in parallel.

But this model had many disadvantages, the most important being the inflexibility associated with it. The queue had a linked list of functions, which could be called, and these would be executed during the process as per the schedule. And the driver could register the bottom halves in respective queues. But this couldn't replace the old BH entirely. Unfortunately, we were unable to handle sub-systems like networking. But during 2.3 kernel development series, programmers addressed this problem by introducing *softirqs* and *tasklets*. The following code illustrates *softirq* 'linked portion' of the interrupt (header) file in kernel source:

```
#define set_softirq_pending(x) (local_softirq_pending() = (x))
#define or_softirq_pending(x) (local_softirq_pending() |= (x))
#endif

/* Some architectures might implement lazy enabling/disabling of
 * interrupts. In some cases, such as stop_machine, we might want
 * to ensure that after a local_irq_disable(), interrupts have
 * really been disabled in hardware. Such architectures need to
 * implement the following hook.
 */
#ifndef hard_irq_disable
#define hard_irq_disable() do { } while(0)
#endif

/* PLEASE, avoid to allocate new softirqs, if you need not _really_ high
frequency threaded job scheduling. For almost all the purposes
tasklets are more than enough. F.e. all serial device BHs et
al. should be converted to tasklets, not to softirqs.
```

```
*/

enum
{
    HI_SOFTIRQ=0,
    TIMER_SOFTIRQ,
    NET_TX_SOFTIRQ,
    NET_RX_SOFTIRQ,
    BLOCK_SOFTIRQ,
    TASKLET_SOFTIRQ,
    SCHED_SOFTIRQ,
#ifdef CONFIG_HIGH_RES_TIMERS
    HRTIMER_SOFTIRQ,
#endif
    RCU_SOFTIRQ, /* Preferable RCU should always be the last softirq
*/

    NR_SOFTIRQS
};

/* softirq mask and active fields moved to irq_cpustat_t in
 * asm/hardirq.h to get better cache usage. KAO
 */

struct softirq_action
{
    void(*action)(struct softirq_action *);
};

asmlinkage void do_softirq(void);
asmlinkage void __do_softirq(void);
extern void open_softirq(int nr, void (*action)(struct softirq_action *));
extern void softirq_init(void);
#define __raise_softirq_irqoff(nr) do { or_softirq_pending(1UL << (nr)); }
while (0)
extern void raise_softirq_irqoff(unsigned int nr);
extern void raise_softirq(unsigned int nr);

/* This is the worklist that queues up per-cpu softirq work.
 *
 * send_remote_sendirq() adds work to these lists, and
 * the softirq handler itself dequeues from them. The queues
 * are protected by disabling local cpu interrupts and they must
 * only be accessed by the local cpu that they are for.
 */
DECLARE_PER_CPU(struct list_head [NR_SOFTIRQS], softirq_work_list);

/* Try to send a softirq to a remote cpu. If this cannot be done, the
 * work will be queued to the local cpu.
 */
extern void send_remote_softirq(struct call_single_data *cp, int cpu, int
softirq);

/* Like send_remote_softirq(), but the caller must disable local cpu
interrupts
 * and compute the current cpu, passed in as 'this_cpu'.
```

```
*/
extern void __send_remote_softirq(struct call_single_data *cp, int cpu,
                                int this_cpu, int softirq);
```

After their introduction the only issue was the compatibility with existing drivers.

softirqs are essentially a set of 32 statically defined bottom halves. They can be made to run in two similar processors, simultaneously. *tasklets*, unlike *softirqs*, are dynamic and are actually built on top of *softirqs*. And they can be made to run on different processors, in parallel.

When you begin actual programming, you can see that *tasklets* are enough for handling your bottom half processing requirements. You may also note that sometimes we need *softirqs* for tasks like networking (owing to reasons concerning performance). The only point you need to keep in your mind is that this requires much attention. Now let's have a glance at the initiation:

```
#ifndef __ARCH_IRQ_STAT
irq_cpustat_t irq_stat[NR_CPUS] ____cacheline_aligned;
EXPORT_SYMBOL(irq_stat);
#endif

static struct softirq_action softirq_vec[32] ____cacheline_aligned_in_smp;
static DEFINE_PER_CPU(struct task_struct *, ksoftirqd);

static inline void wakeup_softirqd(void)
{
    /* Interrupts are disabled; no need to stop preemption */
    struct task_struct *tsk = __get_cpu_var(ksoftirqd);

    if (tsk && tsk->state != TASK_RUNNING)
        wake_up_process(tsk);
}

#define MAX_SOFTIRQ_RESTART 10

asmlinkage void __do_softirq(void)
{
    struct softirq_action *h;
    __u32 pending;
    int max_restart = MAX_SOFTIRQ_RESTART;
    int cpu;

    pending = local_softirq_pending();

    local_bh_disable();
    cpu = smp_processor_id();

restart:
    /* Reset the pending bitmask before enabling irqs */
    local_softirq_pending() = 0;

    local_irq_enable();

    h = softirq_vec;
```

```
do {
    if (pending & 1) {
        h->action(h);
        rcu_bh_qsctr_inc(cpu);
    }
    h++;
    pending >>= 1;
} while (pending);

local_irq_disable();

pending = local_softirq_pending();
if (pending && --max_restart)
    goto restart;

if (pending)
    wakeup_softirqd();

__local_bh_enable();
}

#ifdef __ARCH_HAS_DO_SOFTIRQ

asmlinkage void do_softirq(void)
{
    __u32 pending;
    unsigned long flags;

    if (in_interrupt())
        return;

    local_irq_save(flags);

    pending = local_softirq_pending();

    if (pending)
        __do_softirq();

    local_irq_restore(flags);
}

EXPORT_SYMBOL(do_softirq);

#endif
```

Another aspect that needs the attention of the programmer is that *softirqs* need to be registered statically (during compilation) while the code can dynamically register *tasklets*. You may also note that converting BHs to *softirqs* (or even *tasklets*) is a non-trivial thing!

Fortunately, the 'conversion' later materialised in the 2.5 series development. *tasklet* finally appeared in the apparel of a modified *softirq*, which could be handled easily. (Now you can understand why some authors of *literature-type-texts* refer bottom halves as software interrupts or *softirqs*.) This finally led to

the three bottom-half mechanisms that we mostly deal with (in 2.6 series). The mechanisms are *softirqs*, *tasklets*, and *work queues*.

It is worth mentioning the *kernel timer* here itself. This mechanism essentially performs a postponement of the tasks for specific intervals of time. We will discuss its technical details in the coming days.

In this context, you may have a glance at the code portion that handles *kernel internal timers*, *kernel timekeeping* and *basic process system calls*. This is initiated by:

```
static DEFINE_PER_CPU(tvec_base_t, tvec_bases) = { SPIN_LOCK_
UNLOCKED };

static void check_timer_failed(struct timer_list *timer)
{
    static int whine_count;
    if (whine_count < 16) {
        whine_count++;
        printk("Uninitialised timer!\n");
        printk("This is just a warning. Your computer is OK\n");
        printk("function=0x%p, data=0x%lx\n",
            timer->function, timer->data);
        dump_stack();
    }
    /*
     * Now fix it up
     */
    spin_lock_init(&timer->lock);
    timer->magic = TIMER_MAGIC;
}
```

And the 'starting' is done using:

```
void add_timer_on(struct timer_list *timer, int cpu)
{
    tvec_base_t *base = &per_cpu(tvec_bases, cpu);
    unsigned long flags;

    BUG_ON(timer_pending(timer) || !timer->function);

    check_timer(timer);

    spin_lock_irqsave(&base->lock, flags);
    internal_add_timer(base, timer);
    timer->base = base;
    spin_unlock_irqrestore(&base->lock, flags);
}
```

You can see the actual code that governs *softirqs* in *kernel/softirq.c*. (As I said, we rarely use *softirqs*. In most cases *tasklets* are employed and most of the drivers use *tasklets* for bottom half.) *softirqs* are represented using *softirq_action* which is defined as:

```
struct softirq_action {
    void (*action)(struct softirq_action *); /* function to run */
    void *our_data; /* data to pass to the function */
};
```

Correspondingly, a 32-entry array of this structure can be found in the above code file (*softirq.c*). Since one *softirq* needs one entry, there can be a maximum of 32 registered *softirqs* only. Also, you may see that the kernel actually uses only a fewer entries (out of this 32).

Here is a *softirq* handler for your reference:

```
void softirq_handler(struct softirq_action *)
```

The kernel uses a similar action function with a pointer to the respective *softirq_action* structure, when it runs the *softirq* handler. It is worth mentioning that the kernel passes the entire structure and this facilitates future additions to the structure without redoing the handler. The handler retrieves the data value by dereferencing the argument and looking for the data member. Also, you may find that a *softirq* never attempts to preempt another *softirq*, and the only way to preempt a *softirq* is by deploying an interrupt handler.

You need to check that the registered *softirq* is marked properly before executing it (technically termed as *raising the softirq*) and, normally, the handler marks the corresponding *softirq* for execution before returning.

The pending ones are executed in the following cases:

- Return from a hardware interrupt code
- *ksoftirqd* kernel thread
- Codes that look for pending *softirqs*

The execution occurs when *do_softirq()* is called. You can use this just by using the programming logic: if there are any pending ones, perform *do_softirq()* loop. Now let's look at this part of *do_softirq()*:

```
u32 pending = softirq_pending(cpu);

if (pending) {
    struct softirq_action *h = softirq_vec;

    softirq_pending(cpu) = 0;

    do {
        if (pending & 1)
            h->action(h);
        h++;
    }
```

```

    pending >>= 1;
} while (pending);
}

```

This is the core mechanism associated with its processing. You can see that it looks for pending *softirqs* and executes them. The handler is then registered at run-time via *open_softirq()* which can take three parameters -- *softirq*'s index, handler function and value for the data field.

Now let's come to *tasklets*. First, you need to remember that they have got nothing to do with tasks! As we discussed earlier, *tasklets* also work in a fashion similar to that of *softirqs*. But they have an easy-to-handle interface and extended locking rules. You will be making use of *tasklets* in most cases. Only in very rare cases (when very high-frequency and highly threaded requirements demand it) do we employ *softirqs*. *tasklets* are represented by two *softirqs*: *HI_SOFTIRQ* and *TASKLET_SOFTIRQ* (which are made to run after executing *HI_SOFTIRQ*-based ones). Since they are built on top of *softirqs*, the implementation looks almost similar.

tasklet_struct structure is used to represent the *tasklet* and is defined by:

```

struct tasklet_struct {
    struct tasklet_struct *next; /* next tasklet in the list */
    unsigned long state;        /* state of the tasklet */
    atomic_t count;             /* reference counter */
    void (*func)(unsigned long); /* tasklet handler function */
    unsigned long data;         /* argument to the tasklet function */
};

```

You can employ this in your codes, and the good news is that you don't have to work with the old BH interface any more, as the developers have removed it completely.

Random number generator

Linux has implemented a strong random number generator, which is actually based on the PGP's random number generation method. The generator takes in the 'environmental noise' from device drivers and directs them to an entropy pool. We can access this pool using user and kernel modes. The generator is so effective that the outsider can never predict its value.

Those who know the working mechanism of PGP know the importance of these types of numbers in areas like cryptography. Another point that you need to take into account is that this generator produces true random numbers. They are different from the ones (pseudo-random) that you create using functions in C library—there the problem is that if you know one number in the series, you can guess any other number in the series.

Now let's take a dip in the 'pool' and the associated initiation:

```

static int trickle_thresh = INPUT_POOL_WORDS * 28;

static DEFINE_PER_CPU(int, trickle_count) = 0;

static struct poolinfo {
    int poolwords;
    int tap1, tap2, tap3, tap4, tap5;
} poolinfo_table[] = {
    /* x^128 + x^103 + x^76 + x^51 + x^25 + x + 1 -- 105 */
    { 128, 103, 76, 51, 25, 1 },
    /* x^32 + x^26 + x^20 + x^14 + x^7 + x + 1 -- 15 */
    { 32, 26, 20, 14, 7, 1 },
#ifdef 0
    /* x^2048 + x^1638 + x^1231 + x^819 + x^411 + x + 1 -- 115 */
    { 2048, 1638, 1231, 819, 411, 1 },

    /* x^1024 + x^817 + x^615 + x^412 + x^204 + x + 1 -- 290 */
    { 1024, 817, 615, 412, 204, 1 },

    /* x^1024 + x^819 + x^616 + x^410 + x^207 + x^2 + 1 -- 115 */
    { 1024, 819, 616, 410, 207, 2 },

    /* x^512 + x^411 + x^308 + x^208 + x^104 + x + 1 -- 225 */
    { 512, 411, 308, 208, 104, 1 },

    /* x^512 + x^409 + x^307 + x^206 + x^102 + x^2 + 1 -- 95 */
    { 512, 409, 307, 206, 102, 2 },
    /* x^512 + x^409 + x^309 + x^205 + x^103 + x^2 + 1 -- 95 */
    { 512, 409, 309, 205, 103, 2 },

    /* x^256 + x^205 + x^155 + x^101 + x^52 + x + 1 -- 125 */
    { 256, 205, 155, 101, 52, 1 },

    /* x^128 + x^103 + x^78 + x^51 + x^27 + x^2 + 1 -- 70 */
    { 128, 103, 78, 51, 27, 2 },

    /* x^64 + x^52 + x^39 + x^26 + x^14 + x + 1 -- 15 */
    { 64, 52, 39, 26, 14, 1 },
#endif
};

```

Here the true random number is fully independent of its generating function. Now, let's see how this is done. We know from thermodynamics (Physics!) that entropy is a measurement of disorder and randomness in a system. To represent the randomness in information, John von Neumann suggested the use of the term 'Shannon entropy' and Claude Shannon used it in his theory. Please read the Wikipedia entry on Shannon entropy at [en.wikipedia.org/wiki/Entropy_\(information_theory\)](http://en.wikipedia.org/wiki/Entropy_(information_theory)) for more details.

```

static int debug = 0;

```



```

module_param(debug, bool, 0644);
#define DEBUG_ENT(fmt, arg...) do { if (debug) \
    printk(KERN_DEBUG "random %04d %04d %04d: " \
    fmt, \
    input_pool.entropy_count, \
    blocking_pool.entropy_count, \
    nonblocking_pool.entropy_count, \
    ## arg); } while (0)

#else
#define DEBUG_ENT(fmt, arg...) do {} while (0)
#endif

struct entropy_store;
struct entropy_store {
    /* mostly-read data: */
    struct poolinfo *poolinfo;
    __u32 *pool;
    const char *name;
    int limit;
    struct entropy_store *pull;

    /* read-write data: */
    spinlock_t lock ____cacheline_aligned_in_smp;
    unsigned add_ptr;
    int entropy_count;
    int input_rotate;
};

```

It will be good if you could get a copy of *A Mathematical Theory of Communication* (written by Shannon himself). In the book, the idea of information theory is discussed and Shannon entropy is introduced from scratch.

Shannon entropy is an important concept when we deal with random number generators. Roughly, I can say that high entropy corresponds to 'less useful information', which in turn corresponds to a large amount of random stuff, in a set of characters. Linux has an entropy pool, which has data from non-deterministic device events, making it purely random. It also calculates the 'entropy level change' (technically called *entropy estimate*) when the data is fed into the pool (and is used as a measure of the uncertainty). It does the same when there is a reduction in randomness. *Add_timer_randomness* is shown below:

```

static void add_timer_randomness(struct timer_rand_state *state,
unsigned num)
{
    struct {
        cycles_t cycles;
        long jiffies;
        unsigned num;
    } sample;

```

```

    long delta, delta2, delta3;

extern void add_input_randomness(unsigned int type, unsigned int code,
                                unsigned int value)
{
    static unsigned char last_value;

    /* ignore autorepeat and the like */
    if (value == last_value)
        return;

    DEBUG_ENT("input event\n");
    last_value = value;
    add_timer_randomness(&input_timer_state,
                        (type << 4) ^ code ^ (code >> 4) ^ value);
}

void add_interrupt_randomness(int irq)
{
    if (irq >= NR_IRQS || irq_timer_state[irq] == 0)
        return;

    DEBUG_ENT("irq event %d\n", irq);
    add_timer_randomness(irq_timer_state[irq], 0x100 + irq);
}

void add_disk_randomness(struct gendisk *disk)
{
    if (!disk || !disk->random)
        return;
    /* first major is 1, so we get >= 0x200 here */
    DEBUG_ENT("disk event %d:%d\n", disk->major, disk->first_minor);

    add_timer_randomness(disk->random,
                        0x100 + MKDEV(disk->major, disk->first_minor));
}

EXPORT_SYMBOL(add_disk_randomness);

```

Linux 1.3.30 saw the introduction of kernel random number generator, which is considered as a useful tool by programmers.

With that we have reached the end of today's voyage. Our next destination point will be time management, and later we will move on to kernel synchronisation. Wait till the next instalment to hack more kernel topics.

Happy Kernel Hacking! 

By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V-language—a programming language that employs AI and ANN. His research work/publications are available at www.aasisvinayak.com

Post: Trainee System Administrator

Company: Net Solutions

Profile: Minimum 6 months relevant experience. Candidate should be Microsoft Certified System Engineer (MCSE) or Microsoft Certified Systems Administrator (MCSA). Good knowledge of Microsoft Windows servers, Active Directory, DNS, IIS, DHCP and other Networking Concepts is also desired.
Exp: 0-1, **Location:** Chandigarh
Email: careers@netsolutionsindia.com

Post: Linux Administrator

Company: Technoforte Software Pvt. Ltd.

Profile: Should have good working knowledge on Linux Administration (Preferable on RedHat) & around 2-3 years experience.
Exp: 2-7
Location: Bengaluru
Email: praveen@technoforte.co.in

Post: Sr. Developer- Perl and PHP on Linux/HP-UX

Company: Objectwin Technology India Pvt Ltd

Profile: Should have good communication skills both oral and written with an ability to work independently with the customers elsewhere in the globe. Should have sound working knowledge of quality processes.
Exp: 4-6, **Location:** Bengaluru
Email: nagendra@objectwin.com

Post: System Engineer- Linux

Company: Estel Technologies Pvt. Ltd.
Profile: Should be BE/ B.Tech/ MCA or any equivalent qualification. RHCE certification is must.
Exp: 1-3, **Location:** Pune
Email: jobs@esteltelecom.com

Post: Team Lead- PHP, LINUX

Company: Astrowix India Project Solutions Pvt. Ltd.
Profile: Looking for a candidate who is having 5+ years of development experience of web application in PHP, LINUX Development experience IN Postgre SQL would be Preferred.
Exp: 5-7, **Location:** Noida
Email: sahmad@astrohr.com

Post: Linux Engineer

Company: Netmagic Solutions Pvt. Ltd.

Profile: Strong hands on experience in Linux OS fundamentals & Apache, MySQL, PHP Apps & Strong knowledge of DNS & Knowledge of any Mail application like Qmail or Postfix is desired.
Exp: 5-7, **Location:** Bengaluru
Email: deepthi.anand@netmagicsolutions.com

Post: Sr. Engineers- Linux, Kernel Development

Company: PTL Software India Pvt. Ltd.

Profile: 6-9 years of experience in developing embedded software for Linux/RTOS based platforms is desired.
Exp: 6-9
Location: Bengaluru
Email: rupa_kumbeswaran@phoenix.com

Post: Weblogic Administrator- Unix/Linux, Weblogic, Apache tomcat

Company: First Indian Corporation Pvt. Ltd.

Profile: 5 or more years of experience in Systems Engineering, systems administration, software engineer and network engineering is desired.
Exp: 4-9
Location: Bengaluru
Email: resumesindia@firstam.com

Post: Linux System Testing Engineer

Company: Mani India technologies Pvt. Ltd.

Profile: Good knowledge on Linux is desired & should be Good in analytics.
Exp: 3-5, **Location:** Tuticorin
Email: jobs@maniindiatech.com

Post: Sr. Linux Engineer

Company: Indusa Infotech Services Pvt. Ltd.

Profile: 5+ years of Linux/ Unix programming experience using C++ and Real Time Transaction Processing is desired. An experience in Java, Web Services, SOAP, n-tier design, Network Programming , TCP/IP protocol etc. is desired.
Exp: 5-10, **Location:** Pune
Email: roshan.rawal@indusa.com

Post: System Administrator- Windows, Linux

Company: Treselle Systems Inc

Profile: Experience with Windows, Linux, UNIX administration technology and email systems is desired. Responsible for Monitoring Apache, MySQL and other servers is desired. Will be also required to write scripts to automate monitoring.
Exp: 2-3
Location: Chennai
Email: trjobs@treselle.com

Post: System Admin- AIX, Linux, Solaris, Citrix, WAS, MQ

Company: Synova Innovative Technologies Pvt. Ltd.

Profile: Skills : - AIX HP, UX Linux, Solaris, Citrix, Intel Windows ,WAS Admin, MQSeries Admin. Must have 4+ years relevant experience in one of the above skills.
Exp: 4-7
Location: Chennai
Email: resumes-naukri@synovaindia.com

Post: Unix/ Linux Administrator

Company: TNS India Ltd.

Profile: Any Graduate with an experience of 3+ years in Network Centre Operations (NOC). Good working experience on Unix (preferably HP- UX), Linux , Solaris Administration are essential.
Exp: 3-6
Location: Hyderabad
Email: orsc.vacancy@tns-global.com

Post: Sr. Embedded Engineer- BSP & Device Drivers Development

Company: KPIT Cummins Infosystems Ltd

Profile: Looking for Device Driver Development like BSP, Linux, WinCE etc.
Exp: 5-10
Location: Bengaluru
Email: ravi.h.kumar@kpitcummins.com

Post: Linux Administrator

Company: Intelligroup Asia Pvt. Ltd.

Profile: Primary Skills: Open Source/Linux plus an Experience of 4+ yrs. Experience on at least few of the following services and applications - squid-cache proxy, Apache web-server, sendmail/postfix, samba, NFS, LDAP, Nagios is desired.
Exp: 4-6
Location: Hyderabad
Email: career@intelligroup.com



**Just one right step can help
you reach the top**

**Log on to naukri.com And take your first
step towards success**

**The centre stage for Western India's
most anticipated ICT event**

Book your space Today

www.etict.com

At ET-ICT 2009 International Trade Fair & Conference, some of the key focus areas are VoIP and IP based Next Generation Networks including 3G and Wimax, VAS, Spectrum Utilisation & Infrastructure Sharing, Technologies for E-governance, Health, Education, Induction of IPTV & Mobile TV, new Telecom Policy Initiatives, Investment & Venture Opportunities.

The powerful visitor profile consists of decision makers and high ranking officials from across the region.

By participating in this event you can reach out to all the Enterprise & SME verticals addressed by The Economic Times.

THE ECONOMIC TIMES
presents

International
Trade Fair
& Conference

12-13-14 November 2009
Mumbai, India



your
enterprise
solutions partner

Supporting
Magazine



For Space
Bookings Contact

ET-ICT Secretariat

Anamika Pal - 09821204530

Phone : +91-22-2660 5550 / 7755

Direct : +91-22-6500 2988

Email : sales@etict.com / info@etict.com

The Complete Magazine On Open Source



Features not to be missed

www.linuxforu.com

Technical Content: Technology content—howtos, tips & tricks, coding—catering to System Admins, Developers and Linux Professionals

FOSS Bytes: Latest information on major technology break-through in the field of Linux & Open Source. Key industry announcements & happenings from the world of Linux & Open Source

Reviews: Detailed review of new products and the compatibility with Linux & Open Source

Success Stories

Stories of individuals & companies who have benefited by the use of Linux & Open Source

Case Studies:

Linux development, migration and adoption stories

Subscribe to

'Linux For You' & Save upto 25%+ Get Assured Gifts!!!

A large, stylized number '5' in blue with a thick pink outline. Below the '5' is a yellow banner with the word 'Years' in bold black text.

(60 Issues)

at Rs **4,500**

and Get

8GB Pendrive
worth Rs 1700

FREE



(24 Issues)

at Rs **2,880**

and Get

4GB Pendrive worth Rs 1000

FREE



(12 Issues)

at Rs **1,080**

and Get

2GB Pendrive
worth Rs 600

FREE



Magazine	Period	Cover Price (Rs.)	Offer Price (Rs.)	Assured Gifts
Linux For You (With CD/DVD)	5 Years (60 issues)	6000 <input type="checkbox"/>	4500 <input type="checkbox"/>	8 GB Pendrive
	3 Years (36 issues)	3600 <input type="checkbox"/>	2880 <input type="checkbox"/>	4 GB Pen Drive
	1 Year (12 issues)	1200 <input type="checkbox"/>	1080 <input type="checkbox"/>	2 GB Pen Drive

ORDER FORM

Offer valid till 30th November, 2009

Name	Company Name	Designation
------	--------------	-------------

Mailing Address	Distt.
-----------------	--------

City	Pin Code	State	Phone
------	----------	-------	-------

Email _____ Subscription No. (for existing subscribers only) _____

I would like to subscribe to the above (✓)marked magazine(s), for the period _____. Please find enclosed a sum of Rs _____ by DD/MO/crossed cheque*bearing the No. _____ dt. _____ in favour of **EFY Enterprises Pvt Ltd, payable at Delhi.** (*Please add Rs 50 on non-metro cheques)

or charge my credit card ☐ VISA ☐ MasterCard Please charge Rs.

against my credit card No. | | | | | | | | | | | | | | | |

Date of Birth / / (dd/mm/yy) Card Expiry Date / (mm/yyyy)

Signature (as on the card)



Send this filled-in form or its photocopy to :

EFY Enterprises Pvt Ltd D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020; Ph: 011-26810601-03; Fax: 011-26817563; e-mail: info@efyvindia.com; website: www.efyvindia.com

Terms:- # These rates are applicable for new subscribers as well as renewal by existing subscribers. # The rates are valid for subscribers within India only. # Please allow 4-6 weeks for processing of your subscription. # Subscription Copies will be delivered by ordinary post as soon as the issue is released # Gifts mentioned/Displayed and the gifts sent may be slightly different but will be approximately of the same value # Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.

FOSS Yellow Pages


The best place for you to buy and sell FOSS products and services

HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues

FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories



TARIFF

Category Listing

ONE Category	Rs 2,000
TWO Categories.....	Rs 3,500
THREE Categories.....	Rs 4,750
ADDITIONAL Category	Rs 1,000

Value-add Options

LOGO-plus-Entry.....	Rs 500
Highlight Entry (white background).....	Rs 1,000
Per EXTRA word (beyond 30 words).....	Rs 50

KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.

TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.

✂ Tear & Send

ORDER FORM

Tear & Send ✂

Organisation Name (70 characters): _____

Description (30 words): _____

Email: _____ Website: _____

STD Code: _____ Phone: _____ Mobile: _____

Address (will not be published): _____

City/Town: _____ Pin-code: _____

CATEGORIES

- | | | |
|---|--|---|
| <input type="checkbox"/> CONSULTANTS | <input type="checkbox"/> HIGH PERFORMANCE COMPUTING | <input type="checkbox"/> SOFTWARE DEVELOPMENT |
| <input type="checkbox"/> CONSULTANT (FIRM) | <input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS | <input type="checkbox"/> TRAINING FOR PROFESSIONALS |
| <input type="checkbox"/> EMBEDDED SOLUTIONS | <input type="checkbox"/> LINUX-BASED WEB-HOSTING | <input type="checkbox"/> TRAINING FOR CORPORATE |
| <input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS | <input type="checkbox"/> MOBILE SOLUTIONS | <input type="checkbox"/> THIN CLIENT SOLUTIONS |

Please find enclosed a sum of Rs. _____ by DD/ MO//crossed cheque* bearing the No. _____ dt. _____ in favour of

EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs. 50 on non-metro cheque) towards the cost of _____ FOSS Yellow Pages advertisement(s)

or charge my credit card ☐ VISA ☐ Master Card Please charge Rs. _____

against my credit card No. _____

C V V No. _____ (Mandatory)

Date of Birth ____ / ____ / ____ (dd/mm/yy) Card Expiry Date ____ / ____ (mm/yy)



EFY Enterprises Pvt Ltd., D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com; Website: www.efyindia.com

Signature (as on the card)

To Book Your Listing, Call: Dhiraj (Delhi: 09811206582), Somaiah (B'lore: 09986075717)

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Consultant (Firm)

IB Services

Free Installation of GNU/Linux on Laptops and Desktops. Thin client solutions based on Debian and Ubuntu. Laptops and Desktops pre-installed with Debian and Ubuntu. Migration to GNU/Linux. Data Recovery. Navi Mumbai
Kerala
Mobile: 09847446918
Email: ibmanoj@gmail.com
Web : www.ibservices.in

OS3 Infotech

- Silver Solutions Partner for Novell
- High Availability Computing Solutions
- End-to-end Open Source Solutions Provider
- Certified Red Hat Training Partner
- Corporate and Institutional Training

Navi Mumbai
Mobile: 09324113579
Email: info@os3infotech.com
Web: www.os3infotech.com

Taashee Linux Services

100% Support on LINUX, OSS & JBOSS related projects. We specialize in high-availability and high-performance clusters, remote and onsite system management, maintenance services, systems planning, Linux & JBOSS consulting & Support services.
Hyderabad
Mobile: 09392493753, Fax: 040-40131726
Email: manojkummar@taashee.com
Web: www.taashee.com

Torrid Networks Pvt Ltd

Torrid is a leading provider of information technology consulting with focus on Information Security Services and Open Source Solutions.
Noida
Tel: 0120-4545100, Mobile: 09015505583
Fax: 0120-4235064
Email: tarun.kassana@torridnetworks.com
Web: www.torridnetworks.com

Computer (UMPC) For Linux And Windows

Comptek International

World's smallest computer comptek wibrain B1 UMPC with Linux, Touch Screen, 1 GB RAM 60GB, Wi-Fi, Webcam, upto 6 hour battery (opt.), USB Port, max 1600x1200 resolution, screen 4.8", 7.5"x3.25" Size, weight 526 gm.
New Delhi

Mobile: 09968756177, Fax: 011-26187551
Email: comptekdelhi@compteki.com
Web: www.compteki.com
or www.compteki.in

Education & Training

Aptech Limited

IT, Multimedia and Animation Education and Training
Mumbai
Tel: 022-28272300, 66462300
Fax: 022-28272399
Email: customer@aptech.ac.in
Web: www.aptech-education.com,
www.arena-multimedia.com

IT-Campus: Academy of Information Technology

IT training and solution company with over 12 years of experience. - RHCE
• Software Training • Hardware Training
• Multimedia And Animation • Web Designing • Financial Accounting
Kota (Raj.)
Tel: 0744-2503155, Mobile: 09828503155
Fax: 0744-2505105
Email: m_trilok@yahoo.com
Web: www.doeacc4u.com

Mahan Computer Services (I) Limited

Established in 1990, the organization is primarily engaged in Education and Training through its own & Franchise centres in the areas of IT Software, Hardware, Networking, Retail Management and English. The institute also provides customized training for corporates.
New Delhi
Tel: 011-25916832-33
Email: info@mahanindia.com
Web: www.mahanindia.com

Enterprise Comm. Solutions

Aware Consultants

We specialize in building and managing Ubuntu/Debian Linux servers and provide good dependable system administration. We install and maintain in-house corporate servers. We also provide dedicated and shared hosting as well as reliable wireless/hybrid networking.
Bangalore
Tel: 080-26724324
Email: sales@aware.co.in
Web: www.aware.co.in

Cynapse India Private Limited

We are the creators of open source

product cyn.in. cyn.in is a web 2.0 group collaboration software created by Cynapse, that inter-connects your people with each other and their collective knowledge, seamlessly. It combines the capabilities of collaboration tools like wikis, blogs, file repositories, micro blogs, discussions, audio, videos, and other social applications into a seamless platform. cyn.in helps teams to build collaborative knowledge by sharing and discussing various forms of digital content within a secure, unified application that is accessible using a web based interface or a rich desktop client.
Mumbai

Tel: 022-28445858, 28445629
Email: lisa@cynapse.com
Web: www.cynapse.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.
Bangalore
Tel: 080-40890000
Email: start@deeproot.co.in
Web: www.deeproot.in

ESQUBE Communications Solutions Pvt Ltd

Founders of ESQUBE are faculty at the Indian Institute of Science, Bangalore and carry over eight decades of experience and fundamental knowledge in the field of DSP and Telecommunication. ESQUBE plays a dominant role in the creation of IP in the domain of Sensors, Signals and Systems.

Bangalore
Tel: 080-23517063
Email: info@esqube.com
Web: www.esqube.com

Keen & Able Computers Pvt Ltd

Microsoft Outlook compatible open source Enterprise Groupware
Mobile push, Email Syncing of Contacts/Calendar/Tasks with mobiles
• Mail Archival • Mail Auditing • Instant Messaging
New Delhi
Tel: 011-30880046, 30880047
Mobile: 09810477448, 09891074905
Email: info@keenable.com
Web: www.keenable.com

netCORE
The Innovation Company

Netcore Solutions Pvt Ltd

No.1 company for providing Linux Based Enterprise Mailing solution with around 1500+ Customer all over India. Key Solutions:
• Enterprise Mailing and Collaboration Solution • Hosted Email Security • Mail Archiving Solution • Push Mail on Mobile • Clustering Solution
Mumbai
Tel: 022-66628000
Mobile: 09322985222
Email: kalpit@netcore.co.in
Web: www.netcore.co.in

 **redhat.**

Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.
Mumbai
Tel: 022-39878888
Email: marketing-in@redhat.com
Web: www.redhat.in

Hardware & Networking Institute

Xenitis Technolab Pvt Ltd

Xenitis Technolab is the first of its kind, state-of-the-art infrastructure, Hardware, Networking and I.T Security training institution headquartered in Kolkata. Technolab is the training division of Xenitis group of Companies. It is the proud owner of 'Aamar PC', the most popular Desktop brand of Eastern India. These ranges of PC's are sold in the west under the brand name of 'Aamchi PC', in the north as 'Aapna PC' and in the south as 'Namma PC'.
Kolkata
Tel: 033-22893280
Email: srinku@xenitisgroup.com
Web: www.technolabindia.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

IT Infrastructure Solutions

Absolut Info Systems Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi

Tel: +91-11-26494549

Fax: +91-11-4175 1823

Mobile: +91-9873839960

Email: sales@aisplglobal.com

Web: www.aisplglobal.com

Advent Infotech Pvt Ltd

Advent has an experienced technomarketing team with several years of experience in Networking & Telecom business, and is already making difference in market place. ADVENT qualifies more as Value Added Networking Solution Company, we offers much to customers than just Routers, Switches, VOIP, Network Management Software, Wireless Solutions, Media Conversion, etc.

New Delhi

Tel: 46760000, 09311166412

Fax: 011-46760050

Email: marketingsupport@adventelectronics.com

Web: www.adventelectronics.com

Asset Infotech Ltd

We are an IT solution and training company with an experience of 14 years, we are ISO 9001: 2000. We are partners for RedHat, Microsoft, Oracle and all Major software companies. We expertise in legal software ans solutions.

Dehradun

Tel: 0135-2715965, Mobile: 09412052104

Email: piyush@asset.net.in

Web: www.asset.net.in

BakBone Software Inc.

BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions; proven Solaris products; and application-focused Windows offerings that reliably protect MS SQL, Oracle, Exchange, MySQL and other business critical applications.

New Delhi

Tel: 011-42235156

Email: ashish.gupta@bakbone.com

Web: www.bakbone.com

Clover Infotech Private Limited

Clover Infotech is a leading technology services and solutions provider. Our expertise lies in supporting technology products related to Application, Database, Middleware and Infrastructure. We enable our clients to optimize their business through a combination of best industry practices, standard processes and customized client engagement models. Our core services include Technology Consulting, Managed Services and Application Development Services.

Mumbai

Tel: 022-2287 0659, Fax: 022-2288 1318

Mobile: +91 99306 48405

Email: business@cloverinfotech.com

Web: www.cloverinfotech.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.

Bangalore

Tel: 080-40890000

Email: start@deeproot.co.in

Web: www.deeproot.in

Duckback Information Systems Pvt Ltd

A software house in Eastern India. Business partner of Microsoft, Oracle, IBM, Citrix, Adobe, Redhat, Novell, Symantec, McAfee, Computer Associates, Veritas, Sonic Wall

Kolkata

Tel: 033-22835069, 9830048632

Fax: 033-22906152

Email: asis@duckback.com

Web: www.duckback.co.in

HBS System Pvt Ltd

System Integrators & Service Provider. Partner of IBM, DELL, HP, Sun, Microsoft, Redhat, Trend Micro, Symentic Partners of SUN for their new startup E-commerce initiative

Solution Provider on REDHAT, SOLARIS & JAVA

New Delhi

Tel: 011-25767117, 25826801/02/03

Fax: 25861428

Email: amittal@hbsindia.com

IgnitionWorks

IgnitionWorks provides customized IT solutions for small/medium

businesses by leveraging open source technologies with excellent support & affordable costs. Our expertise include custom software development & infrastructure deployment/support for all industries.

Bangalore

Tel: 080-41243227, Mobile: 9611121452

Email: info@ignitionworks.net

Web: www.ignitionworks.net

INGRES®

Business Open Source

Ingres Corporation

Ingres Corporation is a leading provider of open source database software and support services. Ingres powers customer success by reducing costs through highly innovative products that are hallmarks of an open source deployment and uniquely designed for business critical applications. Ingres supports its customers with a vibrant community and world class support, globally. Based in Redwood City, California, Ingres has major development, sales, and support centers throughout the world, and more than 10,000 customers in the United States and internationally.

New Delhi

Tel: 011-40514199, Fax: +91 22 66459537

Email: sales@ingres.com; info@ingres.com

Web: www.ingres.com

Keen & Able Computers Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi-110019

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com

LDS Infotech Pvt Ltd

Is the authorised partner for RedHat Linux, Microsoft, Adobe, Symantec, Oracle, IBM, Corel etc. Software Services Offered: •Collaborative Solutions •Network Architecture •Security Solutions •Disaster Recovery •Software Licensing •Antivirus Solutions.

Mumbai

Tel: 022-26849192

Email: sales@ldsinfotech.com

Web: www.ldsinfotech.com

Pacer Automation Pvt Ltd

Pacer is leading providers of IT Infrastructure Solutions. We are partners of HP, Redhat, Cisco, Vvmare, Microsoft and Symantec. Our core expertise exists in, Consulting, building and Maintaining the Complete IT Infrastructure.

Bangalore

Tel: 080-42823000, Fax: 080-42823003

Email: solutions@pacerautomation.com

Web: www.pacerautomation.com



Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

Mumbai

Tel: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

Srijan Technologies Pvt Ltd

Srijan is an IT consulting company engaged in designing and building web applications, and IT infrastructure systems using open source software.

New Delhi

Tel: 011-26225926, Fax: 011-41608543

Email: business@srijan.in

Web: www.srijan.in

TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

Mumbai

Tel: 022-66338900; Extn. 324

Email: sales@technoinfotech.com

Web: www.technoinfotech.com

Tetra Information Services Pvt Ltd

One of the leading open source providers. Our cost effective business ready solutions caters of all kind of

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

industry verticles.

New Delhi

Tel: 011-46571313, Fax: 011-41620171

Email: sales@tettrain.com

Web: www.tettrain.com

Tux Technologies

Tux Technologies provides consulting and solutions based on Linux and Open Source software. Focus areas include migration, mail servers, virus and spam filtering, clustering, firewalls, proxy servers, VPNs, server optimization.

New Delhi

Tel: 011-27348104, Mobile: 09212098104

Email: info@tuxtechnologies.co.in

Web: www.tuxtechnologies.co.in

Veeras Infotek Private Limited

An organization providing solutions in the domains of Infrastructure Integration, Information Integrity, Business Applications and Professional Services.

Chennai

Tel: 044-42210000, Fax: 28144986

Email: info@veeras.com

Web: www.veeras.com

Linux-Based Web-Hosting

Manas Hosting

Manas-Hosting is a Bangalore-based company that is dedicated in helping small and midsize business companies to reach customers online. We believe that by creating a website, all you have is just web presence; but to get effective traffic on your website, it is equally important to have a well designed one. This is why we provide the best of Web Hosting and Web Designing services. Also, our services are backed with exceptionally good quality and low costs

Bangalore

Tel: 080-42400300

Email: enquiry@manashosting.com

Web: www.manashosting.com

Linux Desktop

Indserve Infotech Pvt Ltd

OpenLx Linux with Kalcutate (Financial Accounting & Inventory on Linux) offers a complete Linux Desktop for SME users. Its affordable (Rs. 500 + tax as special scheme), Friendly (Graphical UserInterface) and Secure (Virus free).

New Delhi

Tel: 011-26014670-71, Fax: 26014672

Email: info@openlx.com

Web: www.openlx.com

Linux Experts

Intaglio Solutions

We are the training and testing partners of RedHat and the first to conduct RHCCS exam in delhi for the first time ever.

New Delhi

Tel: 011-41582917, 45515795

Email: info@intaglio-solutions.com

Web: www.intaglio-solutions.com

Linux Vendor/Distributors

GT Enterprises

Authorized distributors for Red Hat and JBoss range of products. We also represent various OS's Applications and Developer Tools like SUSE, VMWare, Nokia Qt, MySQL, Codeweavers, Ingres, Sybase, Zimbra, Zend-A PHP Company, High Performance Computing Solutions from The Portland Group, Absoft, Pathscale/Qlogic and Intel Compilers, Scalix-Messaging solution on Linux Platform.

Bangalore

Mobile: +91-9845009939, +91-9343861758

Email : sales@gte-india.com

Web: www.gte-india.com

Taurusoft

Contact us for any Linux Distribution at reasonable rates. Members get additional discounts and Free CD/DVDs with each purchase. Visit our website for product and membership details

Mumbai

Mobile: 09869459928, 09892697824

Email: taurusoft@gmail.com

Web: www.taurusoft.netfirms.com

Software Subscriptions

Blue Chip Computers

Available Red Hat Enterprise Linux, Suse Linux Enterprise Server / Desktop, JBoss, Oracle, ARCServe Backup, AntiVirus for Linux, Verisign/Thawte/GeoTrust SSL Certificates and many other original software licenses.

Mumbai

Tel: 022-25001812, Mobile: 09821097238

Email: bluechip@vsnl.com

Web: www.bluechip-india.com

Software Development

Carizen Software (P) Ltd

Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and fax server. Infrastructure.

Chennai

Tel: 044-24958222, 8228, 9296

Email: info@carizen.com

Web: www.carizen.com

DeepRoot Linux Pvt Ltd

Pure & Exclusive Free Software Business. Creators of the deepOfix Mail Server. We provide: airtight solutions, solid support and Freedom We believe in: sharing, compassion and plain action. Backed by full-time hackers. Quick deployment, easy management. Guaranteed.

Bangalore

Tel: 080-40890000

Email: start@deeproot.co.in

Web: www.deeproot.in



InfoAxon Technologies Ltd

InfoAxon designs, develops and supports enterprise solutions stacks leveraging open standards and open source technologies. InfoAxon's focus areas are Business Intelligence, CRM, Content & Knowledge Management and e-Learning.

Noida

Tel: 0120-4350040, Mobile: 09810425760

Email: sales@infoaxon.com

Web: http://opensource.infoaxon.com

Integra Micro Software Services (P) Ltd

Integra focuses on providing professional services for software development and IP generation to customers. Integra has a major practice in offering Telecom Services and works for Telecom companies, Device Manufacturers, Networking companies, Semiconductor and Application development companies across the globe.

Bangalore

Tel: 080-28565801/05, Fax: 080-28565800

Email: tpvarun@integramicro.com

Web: www.integramicroservices.com

iwebtune.com Pvt Ltd

iwebtune.com is your one-stop, total web site support organisation. We provide high-quality website services and web based software support to any kind of websites, irrespective of the domain or the industry segments.

Bangalore

Tel: 080-4115 2929

Email: santosh@iwebtune.com

Web: www.iwebtune.com

Sarvasv Technologies Pvt Ltd

Web OSX Windows Linux Mobile LAMP W3C apps. Agile, BDD, TDD development. On-Demand, SaaS, ERP, CRM, HRM. Open Source integration. Online project management. VCS (git, svn). automated deployment.

Greater Noida

Mobile: 09711472457, 09810270848

Email: sales@sarvasv.in

Website: www.sarvasv.in

Unistal Systems Pvt Ltd

Unistal is pioneer in Data Recovery Software & Services. Also Unistal is national sales & support partner for BitDefender Antivirus products.

New Delhi

Tel: 011-26288583, Fax: 011-26219396

Email: isales@unistal.com

Web: www.unistal.com

Software and Web Development

Bean eArchitect Integrated Services Pvt Ltd

Application Development, Web Design, SEO, Web Marketing, Web Development.

Navi Mumbai

Tel: 022-27821617, Mobile: 9820156561

Fax: 022-27821617

Email: infodesk@beanarchitect.com

Web: www.beanarchitect.com

Mr Site Takeaway Website Pvt Ltd

Our product is a unique concept in India using which a person without having any technical knowledge can create his website within 1 hour; we also have a Customer Care Center in India for any kind of after sales help. We are already selling it world over with over 65,000 copies sold. It comes with FREE Domain Name, Web Hosting and

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

Customer Care Center for Free Support via Phone and Email and features like PayPal Shopping Cart, Guestbook, Photo Gallery, Contact Form, Forums, Blogs and many more. The price of complete package is just Rs 2,999 per year.

Patiala

Mobile: 91-9780531682

Email: pardeep@mrsite.co.in

Web: www.mrsite.co.in

Salah Software

We are specialized in developing custom strategic software solutions using our solid foundation on focused industry domains and technologies. Also providing superior Solution Edge to our Clients to enable them to gain a competitive edge and maximize their Return on Investments (ROI).

New Delhi

Tel: 011-41648668, 66091565

Email: ceo@salahsoftware.com

Web: www.salahsoftware.com

Thin Client Solutions

Digital Waves

The 'System Integration' business unit offers end-to-end Solutions on Desktops, Servers, Workstations, HPC Clusters, Render Farms, Networking, Security/Surveillance & Enterprise Storage. With our own POWER-X branded range of Products, we offer complete Solutions for Animation, HPC Clusters, Storage & Thin-Client Computing

Mobile: 09880715253

Email: ranga@digitalwaves.in

Web: www.digitalwaves.in

Enjoy Network Solutions

Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment. Active Dealer Channel all over India.

Gujarat

Tel: 0260-3203400, 3241732, 3251732,

Mobile: 09377107650, 09898007650

Email: info@enjoyworld.com

Web: www.enjoyworld.com

Training for Corporate

Bascom Bridge

Bascom Bridge is Red Hat Certified partner for Enterprise Linux 5 and also

providing training to the individuals and corporate on other open source technologies like PHP, MySQL etc.

Ahmedabad

Tel: 079-27545455—66

Fax: 079-27545488

Email: info@bascombridge.com

Web: www.bascombridge.com

Brainnet

Kolkata

Tel: 033-40076450

Email: brainnet@brainware-india.com

Web: www.brainware-india.com

Centre for Excellence in Telecom Technology and Management (CETTM), MTNL

MTNL's Centre for Excellence in Telecom Technology and Management (CETTM) is a state of the art facility to impart Technical, Managerial and corporate training to Telecom; Management personnel. CETTM has AC lecture halls, computer Labs and residential facility.

Mumbai

Tel: 022-25714500, 25714586, 25714585, 25714586

Fax: 022-25706700

Email: contact@cettm.mtnl.in

Web: http://cettm.mtnl.in/infra

Complete Open Source Solutions

RHCT, RHCE and RHCSS training.

Hyderabad

Tel: 040-66773365, 9849742065

Email: nayak.sujeet@gmail.com

Web: www.cossindia.com

ElectroMech

Redhat Linux and open source solution , RHCE, RHCS training and exam center, Ahmedabad and Vadodara

Ahmedabad

Tel: 079-40027898

Email: electromech@electromech.info

Web: www.electromech.info

Focuz Infotech

Focuz Infotech Advanced Education is the quality symbol of high-end Advanced Technology Education in the state. We are providing excellent services on Linux Technology Training, Certifications and live projects to students and corporates, since 2000.

Cochin

Tel: 0484-2335324

Email: enquiry@focuzinfotech.com

Web: www.focuzinfotech.com

G-TEC Computer Education

ISO 9001:2000 certified IT Company, International Testing Centre, Specialised in Multimedia & Animation, conduct MCP, MCSE 2000, MCDBA and MCSA certificates, CCNA, CCNP, the Only authorized centre by INTERNATIONAL AND EUROPEAN COMPUTER UNION to conduct ICDL, Adobe Certifications, training on Web Designing, Tally, Spoken English. Conducts Corporate and institutional training. International certifications issued.

Bangalore

Tel: 080-43567000

Email: gtec.indiranagar@gmail.com

Gujarat Infotech Ltd

GIL is a IT company and 17 years of experience in computer training field. We have experience and certified faculty for the open Source courses like Redhat, Ubuntu, and PHP, MySQL

Ahmedabad

Tel: 079-27452276, Fax: 27414250

Email: info@gujaratinfotech.com

Web: www.gujaratinfotech.com

Lynus Academy Pvt Ltd

India's premier Linux and OSS training institute.

Chennai

Tel: 044-42171278, 9840880558

Email: contactus@lynusacademy.com

Web: www.lynusacademy.com

Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

Bangalore

Tel: 080-22428538, 26600839

Email: info@linuxlearningcentre.com

Web: www.linuxlearningcentre.com

Maze Net Solutions (P) Ltd

Maze Net Solution (P) Ltd, is a pioneer in providing solutions through on time, quality deliverables in the fields of BPO, Software and Networking, while providing outstanding training to aspiring IT Professionals and Call Center Executives. Backed by a team of professional workforce and global alliances, our prime objective is to offer the best blend of technologies in the spheres of Information Technology (IT) and Information Technology Enabled Services (ITES).

Chennai

Tel: 044-45582525

Email: info@mazenetsolution.com

Web: www.mazenetsolution.com

Netweb Technologies

Simplified and scalable storage solutions.

Bangalore

Tel: 080-41146565, 32719516

Email: info@netwebindia.com

Web: www.netwebindia.com

New Horizons India Ltd

New Horizons India Ltd, a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global foot print engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning. Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.

New Delhi

Tel: 011-43612400

Email: info@nhindia.com

Web: www.nhindia.com

Network NUTS

India's only Networking Institute by Corporate Trainers. Providing Corporate and Open classes for RHCE / RHCS training and certification. Conducted 250+ Red Hat exams with 95% result in last 9 months. The BEST in APAC.

New Delhi

Tel: 46526980-2

Mobile: 09310024503, 09312411592

Email: info@networknuts.net

Web: www.networknuts.net

STG International Ltd

An IT Training and Solution Company. Over an experience of 14 years. We are ISO 9001:2000 Certified. Authorised Training Partners of Red Hat & IBM-CEIS. We cover all Software Trainings.

New Delhi

Tel: 011-40560941-42, Mobile:

09873108801

Email: rakhi@stg.in

Web: www.stgonline.com

www.stgglobal.com

TNS Institute of Information Technology Pvt Ltd

Join Red-Hat training and get 100% job guarantee. World's most respected Linux certification. After Red-Hat training, you are ready to join as a Linux Administrator or Network Engineer.

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

New Delhi

Tel: 011-30851100, Fax: 30851103

Email: nks@tiit.co.in

Web: www.tiit.co.in

Webel Informatics Ltd

Webel Informatics Ltd (WIL), a Government of West Bengal Undertaking. WIL is Red Hat Training Partner and CISCO Regional Networking Academy. WIL conducts RHCE, RHCSS, CCNA, Hardware and Software courses.

Kolkata

Tel: 033-22833568, Mobile: 09433111110

Email: enquiry@webelinformatics.com

Web: www.webelinformatics.com

Training for Professionals

AEM

AEM is the Best Certified Redhat Training Partner in Eastern India since last 3 years. AEM conducted more than 500 RHCE exams with 95-100% pass rate. Other courses—RHCSS, SCNA, MCSE, CCNA.

Kolkata

Tel: 033-25488736, Mobile: 09830075018

Email: sinhatuhin1@gmail.com

Web: www.aemk.org

Agam Institute of Technology

In Agam Institute of Technology, we provide hardware and networking training since last 10 years. We specialise in open source operating systems like Red Hat Linux since we are their preferred training partners.

Dehradun

Tel: 0135-2673712, Mobile: 09760099050

Web: www.agamtecindia.com

Amritha Institute of Computer Technology

Amritha Technologies provides an extensive training in high end certification programs and Networking Solutions like Redhat Linux, Redhat Security Services, Cisco, Sun Solaris, Cyber Security Program IBM AIX and so on with a strong focus on quality standards and proven technology processes with most profound principles of Love and Selfless Service.

Mobile: 09393733174

Email: aict.hybd@amrita.ac.in

Web: www.amritahyd.org

Centre For Industrial Research and Staff Performance

A Unique Institute catering to the

need for industries as well as Students for trainings on IT, CISCO certification, PLC, VLSI, ACAD, Pneumatics, Behavior Science and Handicraft.

Bhopal

Tel: 0755-2661412, 2661559

Fax: 0755-4220022

Email: crisp@crispindia.com

Web: www.crispindia.com

Center for Open Source Development And Research

Linux, open source & embedded system training institute and development. All trainings provided by experienced experts & administrators only. Quality training (corporate and individual). We expertise in open source solution. Our cost effective business ready solutions caters of all kind of industry verticals.

New Delhi

Mobile: 09312506496

Email: info@cfosdr.com

Web: www.cfosdr.com

Cisconet Infotech (P) Ltd

Authorised Red Hat Study cum Exam Centre. Courses Offered: RHCE, RHCSS, CCNA, MCSE

Kolkata

Tel: 033-25395508, Mobile: 09831705913

Email: info@cisconetinfo.com

Web: www.cisconetinfo.com

CMS Computer Institute

Red Hat Training partner with 3 Red Hat Certified Faculties, Cisco Certified (CCNP) Faculty, 3 Microsoft Certified Faculties having state Of The Art IT Infrastructure Flexible Batch Timings Available..Leading Networking Institute in Marathwada

Aurangabad

Tel: 0240-3299509, 6621775

Email: aurangabad@cmsinstitute.co.in

Web: www.cmsaurangabad.com

Cyber Max Technologies

OSS Solution Provider, Red Hat Training Partners, Oracle, Web, Thin Clients, Networking and Security Consultancy. Also available CCNA and Oracle Training on Linux. Also available Laptops & PCs

Bikaner

Tel: 0151-2202105, Mobile: 09928173269

Email: cmtech.bikaner@gmail.com,

kr.gupta.ashish@gmail.com

Disha Institute

A franchisee of Unisoft Technologies, Providing IT Training & Computer

Hardware & Networking

Dehradun

Tel: 3208054, 09897168902

Email: thedishainstitute@gmail.com

Web: www.unisofttechnologies.com

EON Infotech Limited (TECHNOSchool)

TechnoSchool is the most happening Training Centre for Red Hat (Linux- Open Source) in the Northern Region. We are fully aware of the Industry's requirement as our Consultants are from Linux industry. We are committed to make you a total industry ready individual so that your dreams of a professional career are fulfilled.

Chandigarh

Tel: 0172-5067566-67, 2609849

Fax: 0172-2615465

Email: info@technoschool.net

Web: http://technoschool.net

GT Computer Hardware Engineering College (P) Ltd

Imparting training on Computer Hardware Networking, Mobile Phone Maintenance & International Certifications

Jaipur

Tel: 0141-3213378

Email: franchise_gt@gteducation.net

Web: www.gteducation.net

HCL Career Development Centre Bhopal

As the fountainhead of the most significant pursuit of human mind (IT), HCL strongly believes, "Only a Leader can transform you into a Leader". HCL CDC is a formalization of this experience and credo which has been perfected over three decades.

Bhopal

Tel: 0755-4094852

Email: bhopal@hclcdc.in

Web: www.hclcdc.in

IINZTRIX E Technologies Pvt Ltd

No. 1 Training prvider in this region.

meerut

Tel: 0121-4020111, 4020222

Mobile: 09927666664

Email: jai@iintrix.com

Web: www.iintrix.com

Indian Institute of Job Oriented Training Centre

Ahmedabad

Tel: 079-40072244—2255—2266

Mobile: 09898749595

Email: info@ijit.net

Web: www.ijit.net

Institute of Advance Network Technology (IANT)

•Hardware Engg. •Networking
•Software Engg. •Multimedia Training.

Ahmedabad

Tel: 079-32516577, 26607739

Fax: 079-26607739

Email: contact@iantindia.com

Web: www.iantindia.com

IPCC

Bridging Gap with professionals.

Lucknow

Tel: 0522-3919496

Email: ipccko@yahoo.co.in

Web: www.ipcc.co.in

Koenig Solutions (P) Ltd

A reputed training provider in India. Authorised training partner of Red Hat, Novell and Linux Professional Institute. Offering training for RHCE, RHCSS, CLP, CLE, LPI - 1 & 2.

New Delhi

Mobile: 09910710143, Fax: 011-25886909

Email: info@koenig-solutions.com

Web: www.koenig-solutions.com

NACS/CIT

We are Providing Training of LINUX to Professional & Cooperate.

Meerut

Tel: 0121-2420587, Mobile: 9997526668

Email: Info@nacsglobal.com

Web: www.nacsglobal.com

NACS Infosystems (P) Ltd

NACS is an organization which is providing training for all international certification, and also NACS is the authorized Training Partner of Redhat and also having testing centre of THOMSON PROMETRIC and PEARSON VUE.

Meerut

Tel: 0121-2767756, Fax: 0121-4006551

Mobile: 09897796603

Email: info@nacsglobal.com,

mohit@nacsglobal.com.

Web: www.nacsglobal.com

Netdiox Computing Systems

We are one-of-a-kind center for excellence and finishing school focusing on ground breaking technology development around distributed systems, networks, storage networks, virtualisation and fundamental algorithms optimized for various appliance.

Bangalore

Tel: 080-26640708

Mobile: 09740846885

Email: info@netdiox.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

NetMax-Technologies

Training Partner of RedHat, Cisco
Chandigarh
Tel: 0172-2608351, 3916555
Email: mail.netmax@gmail.com
Web: www.netmaxtech.com

Netxprt institute of Advance Networking

Netxprt Noida is a Leading organization to provide Open Source training on RedHat Linux RHCT and RHCE Training with 30Hrs. extra exam preparation module.

Noida
Tel: 0120-4346847, Mobile: 09268829812
Email: anoop.vyas@netxpriindia.com
Web: www.netxpriindia.com

Netzone Infotech Services Pvt Ltd

Special batches for MCSE, CCNA and RHCE on RHEL 5 with exam prep module on fully equipped labs including IBM servers, 20+ routers and switches etc. Weekend batches are also available.
New Delhi
Tel: 011-46015674, Mobile: 9212114211
Email: info@netzoneindia.net

Neuron IT Solutions

We offer end to end services and support to implement and manage your IT Infrastructure needs. We also offer Consulting services and Training in Advanced Linux Administration.
Chennai
Mobile: 09790964948
Email: enquiry@neuronit.in
Web: www.neuronit.in

Plexus Software Security Systems Pvt Ltd

Plexus, incorporated in January 2003 is successfully emerged as one of the best IT Company for Networking, Messaging & Security Solutions and Security Training. Networking, Messaging & Security solutions is coupled with the expertise of its training; this has put Plexus in the

unique position of deriving synergies between Networking, Messaging & Security Solutions and IT Training.
Chennai
Tel: 044-2433 7355
Email: training@plexus.co.in
Web: www.plexus.co.in

Professional Group of Education

RHCE & RHCSS Certifications
Jabalpur
Tel: 0761-4039376,
Mobile: 09425152831
Email: naidu.vikas@gmail.com

Q-SOFT Systems & Solutions Pvt Ltd

Q-SOFT is in a unique position for providing technical training required to become a Linux Administration under one roof. Since inception, the commitment of Q-SOFT towards training is outstanding. We Train on Sun Solaris, Suse Linux & Redhat Linux.
Bangalore
Tel: 080-26639207, 26544135, 22440507
Mobile: +91 9945 282834
Email: counsellors@qsoftindia.com
Web: www.qsoftindia.com

Software Technology Network

STN is one of the most acknowledged name in Software Development and Training. Apart from providing Software Solutions to various companies, STN is also involved in imparting High-end project based training to students of MCA and B.Tech etc. of various institutes.
Chandigarh
Tel: 0172-5086829
Email: stn2001@rediffmail.com
Web: stntechnologies.com

South Delhi Computer Centre

SDCC is for providing technical training courses (software, hardware, networking, graphics) with career courses like DOEACC "O" and "A" Level and B.Sc(IT), M.Sc(IT), M.Tech(IT)

from KARNATAKA STATE OPEN UNIVERSITY.
New Delhi
Tel: 011-26183327, Fax: 011-26143642
Email: southdelhicomputercentre@gmail.com, southdelhicomputercentre@hotmail.com.
Web: www.itwhizkid.com
www.itwhizkid.org

Ssytms Quest

Making Tomorrow's professionals TODAY
Bangalore
Tel: 080-41301814
Email: directorv@ssystemsquest.com
Web: www.ssytmsquest.com

Trimax FuturePerfect

A Div of Trimax IT Infrastructure and Services Limited. Redhat RHCE, RHCT Training & Exam Center, MCTS, MCITP, MCSE 03, CCNA, CCNP, Prometric Center.
Mumbai
Tel: 022-40681313, Mobile: 09987705638
Fax: 022-40681001
Email: futureperfect@trimax.in
Web: www.trimax.in

Vibrant e Technologies Ltd

Vibrant e Technologies Ltd. Is a authorised Red Hat Test and Testing Centre, has won the prestigious award " REDHAT BEST CERTIFIED TRAINING PARTNER 2007-2008" for Western region. Vibrant offers courses for RHCE 5, RHCSS etc.
Mumbai
Tel: 022-26285066/6701
Email: vibrant@vsnl.net
Web: www.vibrantcomputers.com

Ultramax Infonet Technologies Pvt Ltd

Training in IT related courses adn authorised testing center of Prometric, Vue and Red Hat.
Mumbai
Tel: 022-67669217
Email: unmesh.raote@ultramaxit.com
Web: www.ultramaxit.com

Yash Infotech

Authorized Training & Exam Center. Best Performing Center in Lucknow for RH Training and Examinations. LINUX & Open Source training institute for IT professionals & Corporate Offering Quality Training for RHCE, RHCSS, PHP, Shell Script, Virtualization and Troubleshooting Techniques & Tools.
Lucknow
Tel: 0522-4043386, Fax: 0522-4043386
Email: yashinfotech.lko@gmail.com

Web Hosting

IDS Logic Pvt Ltd

IDS Logic is a leading innovative IT Solutions company provides Linux Web Hosting, Windows Web Hosting, Application Hosting, Reseller Hosting, Dedicated Server, Virtual Private Server, Server Management & Monitoring, Domain Name Registration, SSL Certificate, Corporate Mailing Solutions, Web Security Solutions.
Noida
Tel: 0120-4235665, Fax: 0120-4235665
Email: info@idslogic.com
Web: www.idswebhosting.com

Perfect Innovation

Web Hosting Spider
• Web Hosting • Web Design
• Web Application Development
• SMS Hosting • Corporate Hosting
• Dedicated Servers
Puducherry
Tel: 0413-3202726, 3246999
Email: admin@webhostingspider.com
Web: www.webhostingspider.com

To advertise in this section,
please contact
Somaiah (B'lore):
09986075717 Dhiraj
(Delhi: 09811206582) on
011-2681-0602 Extn. 222

Want to register your organisation in
FOSS Yellow Pages For **FREE***

Call: Dhiraj (Delhi) 09811206582 Somaiah (Bangalore) 09986075717
or mail: dhiraj.khare@efyindia.com, somaiah.km@efyindia.com

*Offer for limited period.

Surveillance & Security **ZONE**

Securing your Business



South Asia's largest ICT event

18th
Convergence
India 2010

International Exhibition & Conference
Pragati Maidan, New Delhi, India
23-25 March 2010

Supports



Department of Telecommunications
Ministry of Communications & Information Technology
Government of India



Department of Information Technology
Ministry of Communications & Information Technology
Government of India

Supporting Journal



Certified by



Organiser



Exhibitions India Pvt. Ltd.
(An ISO 9001:2000 Certified Company)



Vipul Mehra, General Manager, vipulm@eigroup.in
Tel: +91 11 4279 5000 / 5030 M: +91 99102 04131

www.convergenceindia.org



Cloud Computing

Clickjacking

Data Theft

IP Surveillance

Virtualisation

Info Security



MOBLIN™ ZONE

Create rich “PC-like” internet and media experience on mobile devices with Moblin™ v2

Moblin™ v2 is an optimized open source Linux software stack and application framework for enabling rich Internet experiences on Intel® Atom™ Processor based MIDs, netbooks, nettops, IVI, and embedded systems.

Moblin™ v2 provides the most flexible application framework including a 3D UI for developing Rich Internet and media applications that run across multiple Intel® Atom™ Processor based devices with minimal UI adaptation. The new Moblin™ v2 SDK includes reference code, documentation, and tools to simplify development and provide faster TTM.

To learn more about Moblin™ v2, visit www.moblinzone.com and register to receive regular updates.